



**PUBLIC**

SQL Anywhere - SQL Remote

Document Version: 17.01.0 – 2021-10-15

# SQL Remote

# Content

- 1 SQL Remote. . . . . 4**
- 1.1 SQL Remote Systems. . . . . 5
  - SQL Remote Components. . . . . 6
  - Typical SQL Remote Setups. . . . . 7
  - SQL Remote Replication Process. . . . . 10
- 1.2 Create a SQL Remote System. . . . . 11
  - Publications and Articles. . . . . 13
  - User Privileges. . . . . 22
  - Subscriptions. . . . . 36
  - Transaction Log-based Replication. . . . . 38
  - Replication Conflicts and Errors. . . . . 46
    - Update Conflicts. . . . . 47
    - Row Not Found Errors. . . . . 55
    - Referential Integrity Errors. . . . . 56
    - Duplicate Primary Key Errors. . . . . 58
  - Row Partitioning Among Remote Databases. . . . . 67
    - Disjoint Data Partitions. . . . . 67
    - Overlap Partitions. . . . . 73
  - Unique Identification Numbers for Remote Databases. . . . . 80
- 1.3 SQL Remote Security Considerations. . . . . 83
- 1.4 Management of SQL Remote Systems. . . . . 84
  - Remote Database Extraction. . . . . 86
  - Remote Database Extraction to a Script File. . . . . 88
  - SQL Remote Message Agent (dbremote). . . . . 95
  - SQL Remote Performance. . . . . 102
  - Guaranteed Message Delivery System. . . . . 114
  - Message Size. . . . . 118
  - SQL Remote Message Systems. . . . . 120
  - SQL Remote System Backups. . . . . 138
  - Recovering a Consolidated Database Manually (Command Line). . . . . 145
  - Recovering a Consolidated Database Automatically (Command Line). . . . . 147
  - Replication Error Reporting and Handling. . . . . 149
  - SQL Remote Security Considerations. . . . . 154
  - Upgrades and Resynchronization. . . . . 155
  - SQL Remote Passthrough Mode. . . . . 157
  - Subscription Resynchronization. . . . . 160

1.5	Tutorial: Creating a SQL Remote System. . . . .	165
	Lesson 1: Creating the Consolidated Database. . . . .	166
	Lesson 2: Granting PUBLISH and REMOTE Privileges at the Consolidated Database. . . . .	169
	Lesson 3: Creating Publications and Subscriptions. . . . .	170
	Lesson 4: Creating a SQL Remote Message Type. . . . .	172
	Lesson 5: Extracting the Remote Database. . . . .	173
	Lesson 6: Sending Data from the Consolidated Database to the Remote Database. . . . .	175
	Lesson 7: Receiving Data at the Remote Database. . . . .	176
	Lesson 8: Sending Data from the Remote Database to the Consolidated Database. . . . .	178
1.6	Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server. . . . .	179
	Lesson 1: Creating the Consolidated Database. . . . .	180
	Lesson 2: Configuring the Consolidated Database to Act as the Message Server. . . . .	183
	Lesson 3: Creating the Remote Database. . . . .	184
	Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases. . . . .	186
	Lesson 5: Finishing up. . . . .	189
1.7	Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server. . . . .	190
	Lesson 1: Creating the Consolidated Database. . . . .	191
	Lesson 2: Creating the Message Server. . . . .	193
	Lesson 3: Creating the Remote Database. . . . .	196
	Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases. . . . .	198
	Lesson 5: Finishing up. . . . .	200
1.8	Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server . . . . .	201
	Lesson 1: Creating the Consolidated Database. . . . .	202
	Lesson 2: Configuring the Relay Server. . . . .	205
	Lesson 3: Configuring the Consolidated Database to Act as the Message Server. . . . .	207
	Lesson 4: Creating the Remote Database. . . . .	209
	Lesson 5: Adding and Replicating Data in the Consolidated and Remote Databases. . . . .	211
	Lesson 6: Cleaning Up. . . . .	213
1.9	SQL Remote Reference. . . . .	214
	SQL Remote Utilities and Options Reference. . . . .	214
	SQL Remote System Tables. . . . .	249
	SQL Remote SQL Statements. . . . .	250

# 1 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.

## In this section:

### [SQL Remote Systems \[page 5\]](#)

SQL Remote is a message-based technology designed for the two-way replication of database transactions between a consolidated database and large numbers of remote databases. Administration and resource requirements at the remote sites are minimal, making SQL Remote well suited to mobile devices.

### [Create a SQL Remote System \[page 11\]](#)

Use the consolidated database to complete all SQL Remote administrative tasks.

### [SQL Remote Security Considerations \[page 83\]](#)

SQL Remote provides a number of features to assist in protecting your data.

### [Management of SQL Remote Systems \[page 84\]](#)

You deploy and administer a SQL Remote system from the consolidated database.

### [Tutorial: Creating a SQL Remote System \[page 165\]](#)

In this tutorial you learn how to set up a SQL Remote replication system.

### [Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server \[page 179\]](#)

In this tutorial you learn how to set up a SQL Remote replication system that uses the HTTP messaging system..

### [Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server \[page 190\]](#)

In this tutorial you learn how to set up a SQL Remote replication system that uses the HTTP messaging system with a separate message server.

### [Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server \[page 201\]](#)

In this tutorial you learn how to set up a SQL Remote replication system that uses the Relay Server to forward HTTP traffic to the consolidated database.

### [SQL Remote Reference \[page 214\]](#)

SQL Remote provides utilities, options, procedures, and statements to work with the software.

## 1.1 SQL Remote Systems

SQL Remote is a message-based technology designed for the two-way replication of database transactions between a consolidated database and large numbers of remote databases. Administration and resource requirements at the remote sites are minimal, making SQL Remote well suited to mobile devices.

SQL Remote provides the following functionality:

### **Multiple subscriber support**

SQL Remote allows occasionally connected users to replicate data between a SQL Anywhere consolidated database and a large number of remote SQL Anywhere databases, typically including many mobile databases.

### **Transaction log-based replication**

SQL Remote uses the transaction log for replication. As a result, only changed data is replicated during an update, ensuring proper transaction atomicity throughout the replication system and maintaining consistency among the databases involved in the replication.

### **Central administration**

SQL Remote is centrally administered at the consolidated database. A company can have a large mobile workforce with many unique databases without maintaining each remote database individually. In addition, SQL Remote operation is invisible to the end user.

### **Economical memory use**

To run efficiently, SQL Remote uses memory economically. This allows you to use SQL Remote on existing remote computers and devices without having to invest in new hardware. Replication is possible to and from remote computers and devices with limited space; only relevant data is replicated from the consolidated database to the remote databases.

### **Multi-platform support**

SQL Remote is supported on several operating systems and message links. SQL Anywhere databases can be copied from one file or operating system to another.

### **In this section:**

#### [SQL Remote Components \[page 6\]](#)

A SQL Remote installation is comprised of a database server, SQL Remote, message system client software, and client applications.

#### [Typical SQL Remote Setups \[page 7\]](#)

There are a number of considerations to take into account when deciding whether to use SQL Remote.

#### [SQL Remote Replication Process \[page 10\]](#)

With SQL Remote, messages are always sent in two directions. The consolidated database sends messages containing publication updates to remote databases, and remote databases send updated data and receipt confirmation messages to the consolidated database.

## **Related Information**

[Create a SQL Remote System \[page 11\]](#)

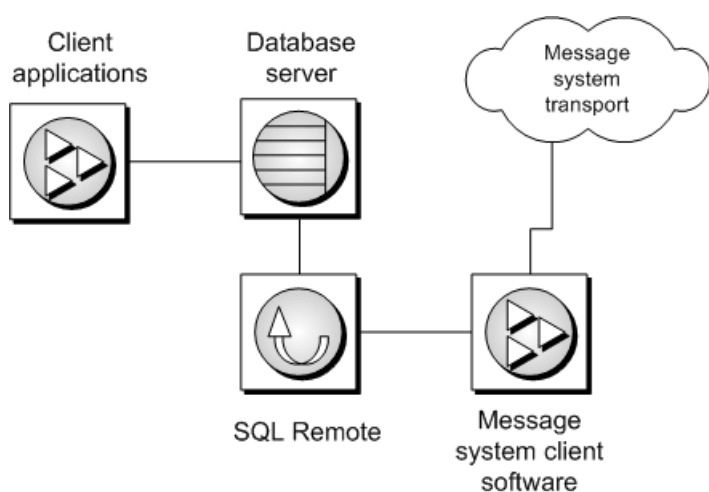
[Supported Platforms](#)

[Management of SQL Remote Systems \[page 84\]](#)

[SQL Remote Reference \[page 214\]](#)

## 1.1.1 SQL Remote Components

A SQL Remote installation is comprised of a database server, SQL Remote, message system client software, and client applications.



### Database server

A SQL Anywhere database is required at the consolidated site and each remote site.

### SQL Remote

To send and receive replication messages from database to database, SQL Remote must be installed at the consolidated site and the remote sites.

The SQL Remote Message Agent connects to the database server via a client/server connection. The SQL Remote Message Agent may run on the same computer as the database server or on a different computer.

### Message system client software

SQL Remote uses existing message systems to transport replication messages.

If you are using a shared file or FTP message system, the message system is included with your operating system.

If you are using an SMTP email system, you must have an email client installed at the consolidated site and each remote site.

### Client applications

The client application can use ODBC, Embedded SQL, or several other programming interfaces. Client applications do not have to know if they are using a consolidated or remote database. From the client application perspective, there is no difference. For specific details about the SQL Anywhere programming interfaces, see the list below.

## Related Information

[JDBC Support](#)  
[ODBC Support](#)  
[OLE DB and ADO Development](#)  
[Perl DBI Support](#)  
[Python and Database Access](#)  
[Ruby Programming](#)  
[SAP Open Client Support](#)  
[SQL Anywhere .NET Data Provider](#)  
[SQL Anywhere C API Support](#)  
[SQL Anywhere PHP Extension](#)  
[Embedded SQL](#)

### 1.1.2 Typical SQL Remote Setups

There are a number of considerations to take into account when deciding whether to use SQL Remote.

SQL Remote is designed for replication systems with the following requirements:

#### **Large numbers of remote databases**

SQL Remote can support thousands of remote databases in a single installation because the messages for many remote databases can be prepared simultaneously.

#### **Occasionally connected**

SQL Remote supports databases that are occasionally connected or indirectly connected to the network. SQL Remote is not designed for instantaneous data availability at each site. For example, it may use an SMTP email system to carry the replication.

#### **Low to high latency**

High latency means a long lag time between data being entered at one database and being replicated to each database in the system. With SQL Remote, replication messages can be sent at intervals of seconds, minutes, hours, or days.

#### **Low to moderate volume**

As replication messages are delivered occasionally, a high transaction volume at each remote database can lead to a large volume of messages. SQL Remote is best suited to systems with a relatively low volume of replicated data per remote database. At the consolidated database, SQL Remote can prepare messages for multiple databases simultaneously.

#### **Homogeneous databases**

Each SQL Anywhere database in the system must have a similar schema.

#### **In this section:**

[Server-to-remote Database Replication for Mobile Workforces \[page 8\]](#)

In this example, SQL Remote provides two-way replication between a consolidated database on an office network, and personal databases on the laptop computers of sales representatives. An SMTP email system is used as a message transport.

## [Server-to-server Database Replication Among Offices \[page 9\]](#)


In this example, SQL Remote provides two-way replication between the database servers at the sales offices or outlets, and the central company office.

### Related Information

[Overview of Data Exchange Technologies](#)  
[How to Choose a Synchronization Technology](#)

## 1.1.2.1 Server-to-remote Database Replication for Mobile Workforces

In this example, SQL Remote provides two-way replication between a consolidated database on an office network, and personal databases on the laptop computers of sales representatives. An SMTP email system is used as a message transport.



To manage the consolidated database, the office network server runs a SQL Anywhere database server. SQL Remote connects to the consolidated database in the same way as any other client application.

Each sales representative's laptop computer includes a SQL Anywhere personal server, a SQL Anywhere remote database, and SQL Remote.

While away from the office, a sales representative can connect to the Internet to run SQL Remote, which carries out the following functions:

- Collects publication updates from the consolidated database on the office network server.



- Submits any local updates, such as new orders, to the consolidated database on the office network server.

The publication updates from the office network database may include new specials on the products the sales representative handles, or new pricing and inventory information. These updates are read by SQL Remote on the laptop and applied to the sales representative's remote database automatically, without requiring any additional action from the sales representative.

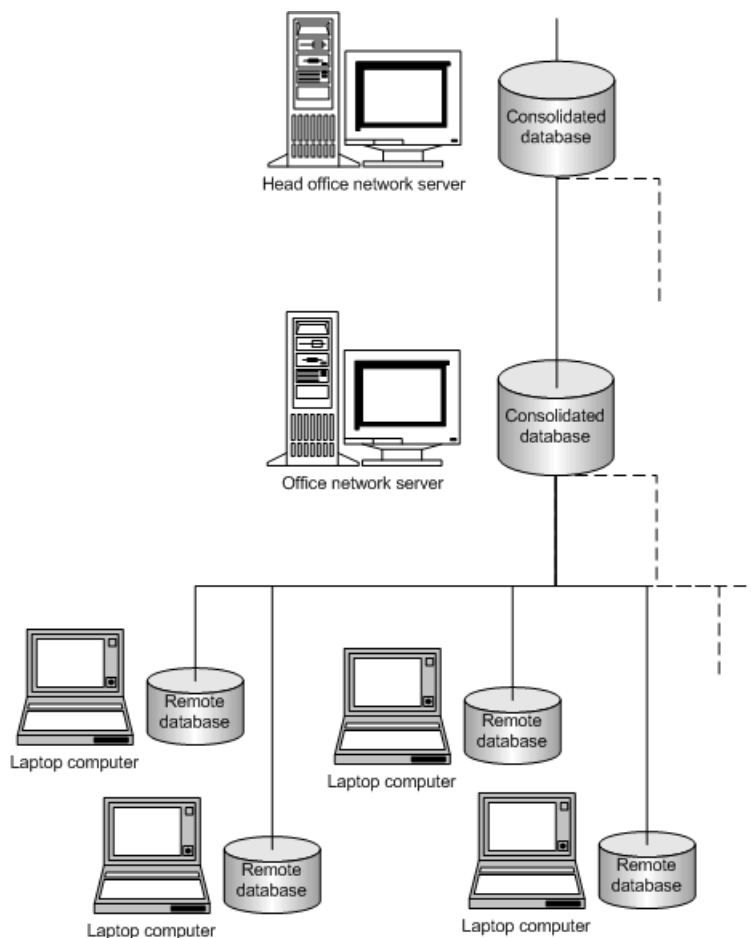
The new orders recorded by the sales representative are also automatically submitted and applied to the office network database without any extra action from the sales representative.

## 1.1.2.2 Server-to-server Database Replication Among Offices

In this example, SQL Remote provides two-way replication between the database servers at the sales offices or outlets, and the central company office.

The only work required at the sales offices is the initial setup and ongoing maintenance of the server.

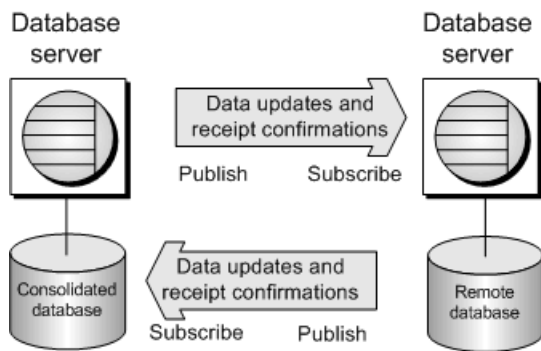
Layers can be added to SQL Remote hierarchies: for example, each sales office server could act as a consolidated database, supporting remote subscribers who work from that office.



SQL Remote can be configured to allow each office to receive its own set of data. Tables such as staff records can be kept private in the same database as the replicated data.

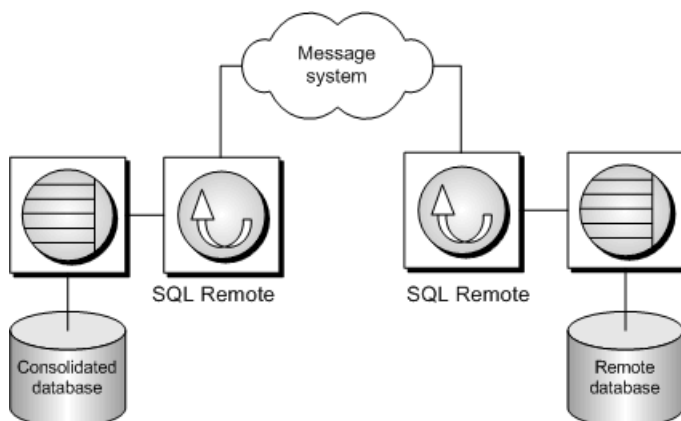
### 1.1.3 SQL Remote Replication Process

With SQL Remote, messages are always sent in two directions. The consolidated database sends messages containing publication updates to remote databases, and remote databases send updated data and receipt confirmation messages to the consolidated database.



When remote database users modify data, their changes are replicated to the consolidated database. When these changes are applied at the consolidated database, they become part of the consolidated database's publication, and are included with the updates sent to all remote databases (except the one the update came from). In this way, replication from remote database to remote database takes place via the consolidated database.

For example, if data in a publication at a consolidated database is updated, those updates are sent to the remote databases. Even if the data is never updated at the remote database, confirmation messages are still sent back to the consolidated database to keep track of the status of the replication.



## Steps Involved in the SQL Remote Replication Process

1. At each consolidated and remote database participating in replication, there is a message agent and a transaction log that manages replication. All committed changes are recorded and stored in the transaction log.
2. Periodically, the SQL Remote Message Agent on the consolidated database scans the transaction log and packages all the committed transactions made to each publication (section of data) into messages. The consolidated database's SQL Remote Message Agent then sends the relevant changes to remote users who subscribe to those publications. The SQL Remote Message Agent sends the changes using a messaging system. SQL Remote supports SMTP email systems, FTP, HTTP/S and FILE.
3. The SQL Remote Message Agent at the remote database accepts the messages sent from the consolidated database and sends a confirmation to the consolidated database that the messages have been received. The SQL Remote Message Agent then applies the transactions to the remote database.
4. At any time, a remote user can run the SQL Remote Message Agent to package the transactions made at the remote database into messages and send them back to the consolidated database.
5. The SQL Remote Message Agent at the consolidated site processes the messages from the remote database and applies the transactions to the consolidated database.

## Related Information

[Create a SQL Remote System \[page 11\]](#)

[Management of SQL Remote Systems \[page 84\]](#)

## 1.2 Create a SQL Remote System

Use the consolidated database to complete all SQL Remote administrative tasks.

To connect to databases using SQL Remote, you must have the SYS\_RUN\_REPLICATION\_ROLE system role.

The following summary describes how to create a SQL Remote system.

1. Choose your SQL Anywhere consolidated database or create a new SQL Anywhere database. The remote databases, which are also SQL Anywhere databases, are created from the consolidated database. When creating a new SQL Anywhere database, keep in mind how SQL Remote uses primary keys (there is the potential for duplicate primary keys when remotes replicate to the consolidated database). A good practice is to choose BIGINT with GLOBAL AUTOINCREMENT for the primary key column data type. Determine what data to replicate.
2. To create an efficient replication system, you need to decide on the tables that you want to use, the columns from those tables, and finally the subset of rows to replicate. Only include the information that is needed.
3. Create publications on the consolidated database. SQL Remote uses a publish and subscribe model to ensure that the correct information reaches its intended user. Arrange the data that you want to replicate into publications on the consolidated database.
4. Create a publisher user on the consolidated database (a publisher is a user with the PUBLISH privilege).

5. Create the remote users on the consolidated database (a remote user is used to uniquely identify a remote database). When you create a remote user, you define the message type to use when transporting the data and, optionally, you define how frequently to send the data.
6. Subscribe the remote users to publications by creating subscriptions.
7. Determine how the remote users can use the data. Remote users can always read their data. You can also allow them to update, delete, and insert data.
8. Choose a method for resolving conflicts. Conflicts can occur during replication when your remote users update, delete, or insert data. You must implement methods for resolving conflicts.
9. Deploy the SQL Remote system to create the remote databases and install the appropriate software.

#### In this section:

##### [Publications and Articles \[page 13\]](#)

A **publication** defines the set of data to be replicated. A publication can include data from several database tables.

##### [User Privileges \[page 22\]](#)

SQL Remote uses a consistent system to manage the users who have privileges on remote and consolidated databases.

##### [Subscriptions \[page 36\]](#)

You subscribe a user to a publication by creating a **subscription**. Each database that shares information in a publication must have a subscription to the publication.

##### [Transaction Log-based Replication \[page 38\]](#)

SQL Remote replicates committed changes and changes that modify data belonging to publications.

##### [Replication Conflicts and Errors \[page 46\]](#)

SQL Remote allows databases to be updated at multiple databases. Careful design is required to avoid replication errors, especially when the database has a complicated structure.

##### [Update Conflicts \[page 47\]](#)

Update conflicts cannot occur when data is shared for reading, or when each row (as identified by its primary key) is updated at only one database. Update conflicts only occur when data is updated at more than one database.

##### [Row Not Found Errors \[page 55\]](#)

A user deletes a row (with a given primary key value.) A second user updates or deletes the same row at another site. In this case, the second statement fails, as the row is not found.

##### [Referential Integrity Errors \[page 56\]](#)

The tables in a relational database are often related through foreign key references. As a result, referential integrity constraints ensure that the database remains consistent.

##### [Duplicate Primary Key Errors \[page 58\]](#)

When all users are connected to the same database, there is no problem ensuring that each INSERT statement uses a unique primary key. If a user tries to re-use a primary key, the INSERT statement fails.

##### [Row Partitioning Among Remote Databases \[page 67\]](#)

Each remote database can contain a different subset of the data stored in the consolidated database. You can create your publications and subscriptions so that data is partitioned among remote databases.

##### [Disjoint Data Partitions \[page 67\]](#)

Data partitioning is disjoint when the remote databases do not share data.

[Overlap Partitions \[page 73\]](#)

Data partitioning overlaps when the remote databases share data. For example, sales representatives can share customers.

[Unique Identification Numbers for Remote Databases \[page 80\]](#)

SQL Remote requires that you assign a different identification number to each remote database.

## Related Information

[GRANT ROLE Statement](#)


[Default Resolution for Update Conflicts \[page 48\]](#)

[Management of SQL Remote Systems \[page 84\]](#)

## 1.2.1 Publications and Articles

A **publication** defines the set of data to be replicated. A publication can include data from several database tables.

An **article** refers to a table in a publication. Each article in a publication can consist of the entire table or a subset of the rows and columns in the table.



## Limitations

A publication cannot include views or stored procedures.

## Viewing Publications and Articles (SQL Central)

In SQL Central, publications appear in the *Publications* folder in the left pane. Any articles you create for a publication appear on the *Articles* tab in the right pane when you select a publication.

### In this section:

#### [Creating Publications \(SQL Central\) \[page 14\]](#)

Create publications based on existing tables in the consolidated database that consist of all the columns and rows in a table.

#### [Publishing Only Some Columns in a Table \(SQL Central\) \[page 15\]](#)

Create a publication that contains all of the rows, but only some of the columns, of a table.

#### [Publishing Only Some Rows in a Table \(SQL Central\) \[page 16\]](#)

To create a publication that contains only some of the rows in a table, you must write a search condition that matches only the rows you want to publish.

#### [Altering a Publication \(SQL Central\) \[page 20\]](#)

Alter a publication by adding, modifying, deleting articles, or by renaming the publication.

#### [Dropping a Publication \(SQL Central\) \[page 21\]](#)

Drop a publication (all subscriptions to that publication are automatically deleted).

## Related Information

[Procedure Replication \[page 42\]](#)

[Trigger Replication \[page 42\]](#)

### 1.2.1.1 Creating Publications (SQL Central)

Create publications based on existing tables in the consolidated database that consist of all the columns and rows in a table.

#### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

#### Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the consolidated database.

2. In the left pane, click the *Publications* folder.
3. Click **File > New > Publication**.
4. In the *What Do You Want To Name The New Publication* field, type a name for the publication. Click *Next*.
5. In the *Choose a Publication Type* window, select the appropriate publication type. Click *Next*.
6. In the *Specify Tables* list, click one or more (using Ctrl-click) tables. Click *Add*.
7. Click *Finish*.

## Results

The publication is created as specified.

## Related Information

[CREATE PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

### 1.2.1.2 Publishing Only Some Columns in a Table (SQL Central)

Create a publication that contains all of the rows, but only some of the columns, of a table.

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the consolidated database.
2. In the left pane, expand the *Publications* folder.
3. Click **File > New > Publication**.
4. In the *What Do You Want To Name The New Publication* field, type a name for the publication. Click *Next*.
5. In the *Choose a Publication Type* window, select the appropriate publication type. Click *Next*.
6. In the *Specify Tables* list, click one or more (using Ctrl-click) tables. Click *Add*.
7. On the *Available Columns* tab, click the table's icon to expand the list of *Available Columns*. Click each column you want to publish and click *Add*. Click *Next*.

8. In the *Create Upload Procedures* window, choose which stored procedures to create, as necessary.
9. Click *Finish*.

## Results

The publication is created as specified.

## Next Steps

Create a subscription.

## Related Information

[Subscriptions \[page 36\]](#)

[CREATE PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

[Referential Integrity Errors \[page 56\]](#)

### 1.2.1.3 Publishing Only Some Rows in a Table (SQL Central)

To create a publication that contains only some of the rows in a table, you must write a search condition that matches only the rows you want to publish.

Use one of the following clauses in your search condition:

#### **SUBSCRIBE BY clause**

Use the SUBSCRIBE BY clause when multiple subscribers to a publication receive different rows from a table.

The SUBSCRIBE BY clause is recommended when your SQL Remote system requires a large number of subscriptions. The SUBSCRIBE BY clause allows many subscriptions to be associated with a single publication, whereas the WHERE clause does not. Subscribers receive rows depending on the value of a supplied expression.

Publications using a SUBSCRIBE BY clause are more compact, easier to understand, and provide better performance than maintaining several WHERE clause publications.

#### **WHERE clause**

Use a WHERE clause to include a subset of rows in an article. All subscribers to the publication containing this article receive the rows that satisfy the WHERE clause.

All unpublished rows must have a default value. Otherwise, when the remote database tries to insert new rows from the consolidated database, an error occurs.



You can combine a WHERE clause in an article.

The database server must add information to the transaction log, and scan the transaction log to send messages, in direct proportion to the number of publications. The WHERE clause does not allow many subscriptions to be associated with a single publication; however the SUBSCRIBE BY clause does.

## Example

You need a publication that enables each sales representative to:

- Subscribe to their sales orders.
- Update their sales orders locally.
- Replicate their sales to the consolidated database.

If you use the WHERE clause, you would need to create separate publications for each sales representative. The following publication is for a sales representative named Sam Singer; each of the other sales representatives would need a similar publication.

```
CREATE PUBLICATION PubOrdersSamSinger (  
    TABLE SalesOrders  
    WHERE Active = 1  
);
```

The following statement subscribes Sam Singer to the PubsOrdersSamSinger publication.

```
CREATE SUBSCRIPTION  
TO PubOrdersSamSinger  
FOR Sam_Singer;
```

If you use the SUBSCRIBE BY clause, you need only one publication. All of the sales representatives can use the following publication:

```
CREATE PUBLICATION PubOrders (  
    TABLE SalesOrders  
    SUBSCRIBE BY SalesRepresentativeID  
);
```

The following statement subscribes Sam Singer to the PubsOrders publication by his ID, 8887.

```
CREATE SUBSCRIPTION  
TO PubOrders ('8887')  
FOR Sam_Singer;
```

### In this section:

[Publishing Only Some Rows Using the SUBSCRIBE BY Clause \(SQL Central\) \[page 18\]](#)

Create a publication using the SUBSCRIBE BY clause.

[Publishing Only Some Rows Using a WHERE Clause \(SQL Central\) \[page 19\]](#)

Create a publication that uses a WHERE clause to include all the columns, but only some of the rows of a table.

## Related Information

[CREATE PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

[CREATE SUBSCRIPTION Statement \[SQL Remote\]](#)

### 1.2.1.3.1 Publishing Only Some Rows Using the SUBSCRIBE BY Clause (SQL Central)

Create a publication using the SUBSCRIBE BY clause.

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the consolidated database.
2. In the left pane, double-click *Publications*.
3. Click **File > New > Publication**.
4. In the *What Do You Want To Name The New Publication* field, type a name for the publication. Click *Next*.
5. Click *Next*.
6. In the *Available Tables* list, click a table. Click *Add*. Click *Next*.
7. On the *Available Columns* tab, click the table's icon to expand the list of *Available Columns*. Click each column you want to publish and click *Add* (use Ctrl-click to select multiple columns at a time). Click *Next*.
8. Click *Next*.
9. On the *Specify SUBSCRIBE BY Restrictions* page:
  - a. Click a table in the *Articles* list.
  - b. Click *Column* and click a column from the dropdown list.
10. Click *Finish*.

## Results

The publication is created as specified.

## Next Steps

Subscribe a user to the publication.

## Related Information

[Creating a SQL Remote Subscription \(SQL Central\) \[page 37\]](#)

[Publishing Only Some Rows Using a WHERE Clause \(SQL Central\) \[page 19\]](#)

[Disjoint Data Partitions \[page 67\]](#)

[CREATE PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

[CREATE SUBSCRIPTION Statement \[SQL Remote\]](#)

### 1.2.1.3.2 Publishing Only Some Rows Using a WHERE Clause (SQL Central)

Create a publication that uses a WHERE clause to include all the columns, but only some of the rows of a table.

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Context

Users can subscribe to more than one publication, and can have more than one subscription to a single publication.

## Procedure

1. Using SQL Central, connect to the consolidated database.
2. Double-click the *Publications* folder.
3. Click **File** > **New** > **Publication**.
4. In the *What Do You Want To Name The New Publication* field, type a name for the publication. Click *Next*.
5. Click *Next*.
6. In the *Available Tables* list, click a table. Click *Add*. Click *Next*.

7. On the *Available Columns* tab, double-click the table's icon to expand the list of *Available Columns*. Click each column you want to publish and click *Add*. Click *Next*.
8. On the *Specify WHERE Clauses* page:
  - a. Click a table in the *Articles* list.
  - b. Type a WHERE clause into the *The Selected Article Has The Following WHERE Clause* field.
9. Click *Finish*.

## Results

The publication is created as specified.

## Next Steps

Add a subscription.

## Related Information

[Subscriptions \[page 36\]](#)

[Row Not Found Errors \[page 55\]](#)

[CREATE PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

[CREATE SUBSCRIPTION Statement \[SQL Remote\]](#)

### 1.2.1.4 Altering a Publication (SQL Central)

Alter a publication by adding, modifying, deleting articles, or by renaming the publication.

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Context

### ⚠ Caution

Altering publications in a running SQL Remote system can cause replication errors and a loss of data in the replication system.

## Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click *Publications*.
3. Right-click the publication you want to alter and click *Properties* to edit the publication.

## Results

The publication is altered.

## Related Information

[Upgrades and Resynchronization \[page 155\]](#)

[ALTER PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

### 1.2.1.5 Dropping a Publication (SQL Central)

Drop a publication (all subscriptions to that publication are automatically deleted).

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Context

### ⚠ Caution

Dropping publications in a running SQL Remote system can cause replication errors and a loss of data in the replication system.

## Procedure

1. In SQL Central, use the [SQL Anywhere 17](#) plug-in to connect to the database.
2. Double-click [Publications](#).
3. Right-click the desired publication and click [Delete](#).

## Results

The publication is dropped and all subscriptions to that publication are deleted.

## Related Information

[Upgrades and Resynchronization \[page 155\]](#)

[DROP PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

## 1.2.2 User Privileges

SQL Remote uses a consistent system to manage the users who have privileges on remote and consolidated databases.

Users of databases involved in SQL Remote replication require the `MANAGE REPLICATION` system privilege (part of the `SYS_REPLICATION_ADMIN_ROLE` system role) to `GRANT` or `REVOKE` the following privileges:

### **PUBLISH**

Every database in a SQL Remote system publishes information. Therefore, every database must have a publisher. To create a publisher, grant one user `PUBLISH` privilege. The publisher user must be unique throughout the SQL Remote system. When sending data, the publisher represents the database. For example, when a database sends a message, its publisher user name is included with the message. When a database receives a message, it can identify the database that sent the message by the publisher name in the message.

### **REMOTE**

A database, such as a consolidated database, that sends messages to other databases must specify which remote databases it sends messages to. To specify these remote databases on the consolidated database, grant REMOTE privilege to the publishers of the remote databases. REMOTE privilege identifies databases that receive messages from the current database.

### **CONSOLIDATE**

Each remote database must specify the consolidated database that it receives messages from. To specify a consolidated database on the remote database, grant CONSOLIDATE privilege to the publisher of the consolidated database. A remote database can only receive messages from one consolidated database. CONSOLIDATE privilege identifies the database that sends messages to this remote database.

## **Extraction Utility (dbxtract) Sets Privileges Automatically**

By default, the Extraction utility (dbxtract) and the *Extract Database Wizard* grant the appropriate PUBLISH and CONSOLIDATE privileges to users in the remote databases.

### **In this section:**

#### [Single-tiered Hierarchy \[page 24\]](#)

In a single-tiered hierarchy, there is one consolidated database with one or more remote databases underneath. In such a hierarchy, the consolidated database grants REMOTE privilege to the publishers of the remote databases.

#### [Multi-tiered Hierarchy \[page 25\]](#)

In a multi-tier hierarchy, all remote databases immediately below the current database are granted REMOTE privilege. The database immediately above the current database in the hierarchy is granted CONSOLIDATE privilege.

#### [PUBLISH Privilege \[page 26\]](#)

Every database in a SQL Remote system requires a publisher, which is a unique user with PUBLISH privilege. You create a publisher by granting a user PUBLISH privilege.

#### [REMOTE Privilege \[page 29\]](#)

Granting REMOTE privilege is also referred to as adding a remote user to the database. Publishers of databases directly below the current database in a SQL Remote hierarchy are granted REMOTE privilege by the current database.

#### [CONSOLIDATE Privilege \[page 32\]](#)

Databases directly above the current database in a SQL Remote hierarchy are granted CONSOLIDATE privilege by the current database. At each remote database, the consolidated database must be granted CONSOLIDATE privilege.

#### [Replication-related System Roles for SQL Remote \[page 35\]](#)

There are two replication-related SQL Anywhere system roles that control access and capabilities in SQL Remote: SYS\_RUN\_REPLICATION\_ROLE, and SYS\_REPLICATION\_ADMIN\_ROLE.

## **Related Information**

[Extraction Utility \(dbxtract\) \[page 225\]](#)

## 1.2.2.1 Single-tiered Hierarchy

In a single-tiered hierarchy, there is one consolidated database with one or more remote databases underneath. In such a hierarchy, the consolidated database grants REMOTE privilege to the publishers of the remote databases.

Each remote database grants CONSOLIDATE privilege to the consolidated database publisher.

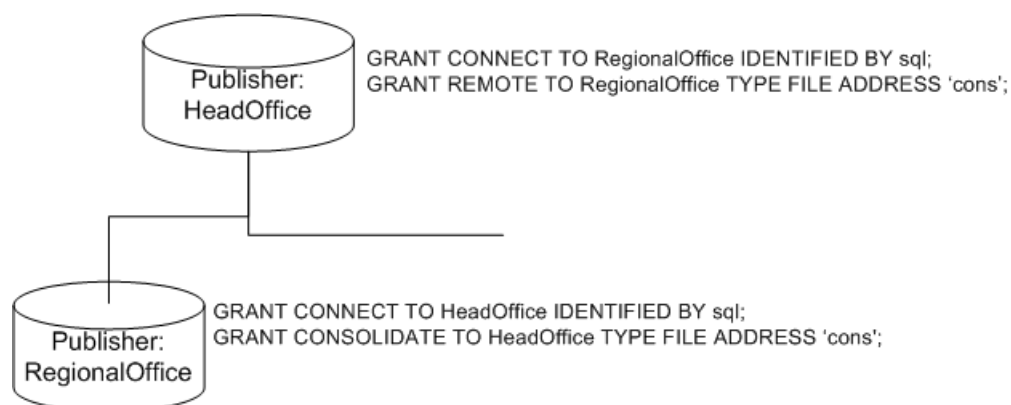
For example, there is a consolidated database identified by its publisher, HeadOffice, and a remote database identified by its publisher, RegionalOffice.

On the consolidated database, HeadOffice, you:

- Create a user with the same name as the publisher of the remote database: RegionalOffice.
- Grant REMOTE privilege to RegionalOffice. This identifies RegionalOffice as a database that receives messages from HeadOffice.

On the remote database, RegionalOffice, you:

- Create a user with the same name as the publisher of the consolidated database: HeadOffice.
- Grant CONSOLIDATE privilege to HeadOffice. This identifies HeadOffice as the consolidated database for RegionalOffice; that is, HeadOffice is the database that sends messages to RegionalOffice.



## Dbxtract Sets Privileges Automatically

By default, the Extraction utility (dbxtract) and the *Extract Database Wizard* grant the appropriate PUBLISH and CONSOLIDATE privileges to users in the remote databases.

## Related Information

[Extraction Utility \(dbxtract\) \[page 225\]](#)



## 1.2.2.2 Multi-tiered Hierarchy

In a multi-tier hierarchy, all remote databases immediately below the current database are granted REMOTE privilege. The database immediately above the current database in the hierarchy is granted CONSOLIDATE privilege.

For example, there is a consolidated database identified by its publisher, HeadOffice, which has a remote database, RegionalOffice. However, the RegionalOffice database also has a remote database: Office.

On the consolidated database, HeadOffice, you:

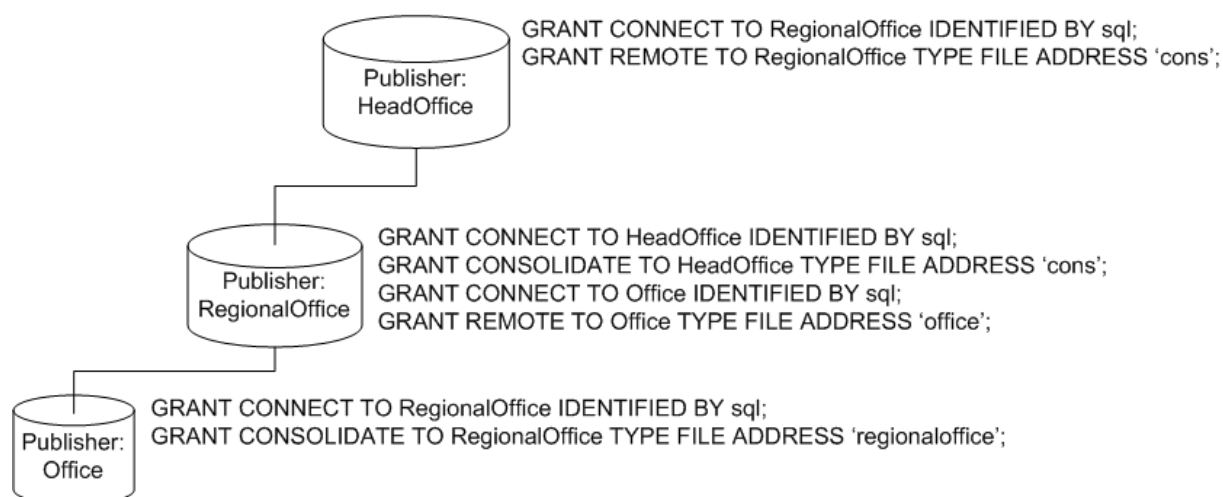
- Create a user with the same name as the publisher of the remote database RegionalOffice.
- Grant REMOTE privilege to the user RegionalOffice. This identifies RegionalOffice as a database that receives messages from HeadOffice.

On the RegionalOffice database, you:

- Create a user with the same name as the publisher of the consolidated database HeadOffice.
- Grant CONSOLIDATE privilege to HeadOffice. This identifies HeadOffice as the consolidated database for RegionalOffice; that is, HeadOffice is the database that sends messages to RegionalOffice.
- Create a user with the same name as the database immediately below RegionalOffice: Office.
- Grant REMOTE privilege to Office. This identifies Office as a database that receives messages from RegionalOffice.

On the Office database, you:

- Create a user with the same name as the publisher of the consolidated database: RegionalOffice.
- Grant Consolidate privilege to the RegionalOffice user. This identifies RegionalOffice as the consolidated database for Office; that is RegionalOffice sends messages to Office.



## 1.2.2.3 PUBLISH Privilege

Every database in a SQL Remote system requires a publisher, which is a unique user with PUBLISH privilege. You create a publisher by granting a user PUBLISH privilege.

All outgoing SQL Remote messages, including publication updates and receipt confirmations, are identified by their publisher. Every database in a SQL Remote system sends receipt confirmations.

The PUBLISH privilege has no authority except to identify the publisher in outgoing messages.

When PUBLISH privilege is granted to a user-extended role, it is not inherited by members of the role.

### In this section:

[Creating a Publisher \(SQL Central\) \[page 26\]](#)

Create users and grant them PUBLISH privilege.

[Revoking PUBLISH Privilege \(SQL Central\) \[page 27\]](#)

Revoke the PUBLISH privilege from a user.

[Viewing the Publisher \(SQL Central\) \[page 28\]](#)

Identify users that are publishers.

### 1.2.2.3.1 Creating a Publisher (SQL Central)

Create users and grant them PUBLISH privilege.

#### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

#### Context

When a remote database is extracted by the Extraction utility (dbxtract) or the *Extract Database Wizard*, the remote user becomes the publisher of the remote database, and is granted PUBLISH privilege.

#### Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click the database name.
3. Double-click *Users*.

4. Right-click the user name.
5. Click *Change to Publisher*.

## Results

The user is created and they have PUBLISH privilege.

## Related Information

[Revoking PUBLISH Privilege \(SQL Central\) \[page 27\]](#)

[Viewing the Publisher \(SQL Central\) \[page 28\]](#)

[GRANT Statement](#)

[GRANT PUBLISH Statement \[SQL Remote\]](#)

### 1.2.2.3.2 Revoking PUBLISH Privilege (SQL Central)

Revoke the PUBLISH privilege from a user.

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Context

### Caution

Changing the publisher at a remote database or at the consolidated database can cause serious problems for any subscriptions that the database is involved in, including the loss of information.

## Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click the database name.

3. Double-click *Users*.
4. Right-click the user name.
5. Click *Revoke Publisher*.

## Results

The user no longer has PUBLISH privilege.

## Related Information

[Viewing the Publisher \(SQL Central\) \[page 28\]](#)

[REVOKE PUBLISH Statement \[SQL Remote\]](#)

[Changes to Avoid on a Running System \[page 156\]](#)

### 1.2.2.3.3 Viewing the Publisher (SQL Central)

Identify users that are publishers.

## Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click *Users*.

The publisher is the user whose *Type* is *Publisher*.

## Results

You are able to view the users that are publishers.

## Related Information

[Revoking PUBLISH Privilege \(SQL Central\) \[page 27\]](#)

[db\\_publisher Option](#)

## 1.2.2.4 REMOTE Privilege

Granting REMOTE privilege is also referred to as adding a remote user to the database. Publishers of databases directly below the current database in a SQL Remote hierarchy are granted REMOTE privilege by the current database.

When granting REMOTE privilege to a user, you must configure the following settings:

### Message system

You cannot create a new remote user until at least one message system is defined in the database.

### Send frequency

When you use SQL statements to grant REMOTE privilege, setting the send frequency is optional.

Granting REMOTE privilege to a user:

- Identifies the user as a remote user.
- Specifies a message type to use for exchanging messages with this remote user.
- Provides an address to send messages to.
- Indicates how often messages should be sent to the remote user.

The publisher for a database cannot have REMOTE and CONSOLIDATE privilege on the same database. This would identify the publisher as both the sender and recipient of outgoing messages.

## Granting REMOTE Privilege to Groups

Although, you can grant REMOTE privilege to a user-extended role, the REMOTE privilege is *not* inherited by grantees of the role. You must explicitly grant REMOTE privilege to each grantee of a user-extended role.

### In this section:

[Granting REMOTE Privilege \(SQL Central\) \[page 30\]](#)

Add a remote user or change an existing user to a remote user.

[Revoking REMOTE Privilege \(SQL Central\) \[page 31\]](#)

Remove a user or role from the SQL Remote system, revert that user or role to a normal user/role, and unsubscribe the user or role from all publications.

## Related Information

[SQL Remote Message Systems \[page 120\]](#)

[Send Frequency \[page 99\]](#)

[GRANT REMOTE Statement \[SQL Remote\]](#)

## 1.2.2.4.1 Granting REMOTE Privilege (SQL Central)

Add a remote user or change an existing user to a remote user.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. In SQL Central, use the [SQL Anywhere 17](#) plug-in to connect to the database.
2. Double-click the database name.
3. Double-click [Users](#).
4. Right-click the user name.
5. Choose one of the following:

Option	Action
<b>Add a new remote user</b>	<ol style="list-style-type: none"><li>1. Click <a href="#">File</a> &gt; <a href="#">New</a> &gt; <a href="#">User</a>.</li><li>2. Follow the instructions in the <a href="#">Create Remote User Wizard</a>.</li><li>3. Right-click the user and click <a href="#">Change to Remote User</a>.</li></ol>
<b>Make an existing user a remote user</b>	<ol style="list-style-type: none"><li>1. Right-click a user and click <a href="#">Change To Remote User</a>.</li><li>2. In the window, click the message type, enter an address, click the send frequency, and click <a href="#">OK</a>.</li></ol>

### Results

The remote user is created.

### Related Information

[Revoking REMOTE Privilege \(SQL Central\) \[page 31\]](#)

[GRANT Statement](#)

[GRANT REMOTE Statement \[SQL Remote\]](#)

[CREATE USER Statement](#)

## 1.2.2.4.2 Revoking REMOTE Privilege (SQL Central)

Remove a user or role from the SQL Remote system, revert that user or role to a normal user/role, and unsubscribe the user or role from all publications.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

Revoking a user or group's REMOTE privilege:

- Removes a user from the SQL Remote system.
- Reverts that user or group to a normal user/group.
- Unsubscribes the user or group from all publications.

### Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click the database name.
3. Double-click *Users*.
4. Right-click the user name.
5. Click *Revoke Remote*.

### Results

The user/group is reverted to a normal user/group, removed from the SQL Remote system, and unsubscribed from all publications.

### Related Information

[Granting REMOTE Privilege \(SQL Central\) \[page 30\]](#)

[REVOKE REMOTE Statement \[SQL Remote\]](#)

## 1.2.2.5 CONSOLIDATE Privilege

Databases directly above the current database in a SQL Remote hierarchy are granted CONSOLIDATE privilege by the current database. At each remote database, the consolidated database must be granted CONSOLIDATE privilege.

CONSOLIDATE privilege must be granted even from read-only remote databases to the consolidated database, because receipt confirmations are sent from the remote databases to the consolidated database.

When granting CONSOLIDATE privilege to a user, you must configure the following settings:

### Message system

You cannot create a new consolidated user until at least one message system is defined in the database.

### Send frequency

When you use SQL statements to grant CONSOLIDATE privilege, setting the send frequency is optional.

Granting CONSOLIDATE privilege:

- Identifies a user as a consolidated user.
- Specifies a message type to use for exchanging messages with this consolidated user.
- Provides an address to send messages to.
- Indicates how often messages should be sent to the consolidated user.

The publisher for a database cannot have REMOTE and CONSOLIDATE privileges on the same database. This would identify the publisher as both the sender and recipient of outgoing messages.

## Extraction utility (dbxtract)

When a remote database is extracted by the Extraction utility (dbxtract) or the *Extract Database Wizard*, the GRANT CONSOLIDATE statement is executed automatically on the remote database.

### In this section:

[Granting CONSOLIDATE Privilege \(SQL Central\) \[page 33\]](#)

Grant CONSOLIDATE privilege to a user.

[Revoking CONSOLIDATE Privilege \(SQL Central\) \[page 34\]](#)

Remove the CONSOLIDATE privilege for a user of the database.

## Related Information

[SQL Remote Message Systems \[page 120\]](#)

[Send Frequency \[page 99\]](#)



## 1.2.2.5.1 Granting CONSOLIDATE Privilege (SQL Central)

Grant CONSOLIDATE privilege to a user.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

Grant CONSOLIDATE privilege to the publisher of the consolidated database.

### Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click the database name.
3. Double-click *Users*.
4. Right-click the user name.
5. Click *Change to Consolidated User*.
6. Configure the *Message Type*, *Address*, and *Send Frequency* settings.
7. Click *OK* to close the *Change to Consolidated User* window.

### Results

The user now has CONSOLIDATE privilege.

### Related Information

[Revoking CONSOLIDATE Privilege \(SQL Central\) \[page 34\]](#)

[GRANT Statement](#)

[GRANT CONSOLIDATE Statement \[SQL Remote\]](#)

## 1.2.2.5.2 Revoking CONSOLIDATE Privilege (SQL Central)

Remove the CONSOLIDATE privilege for a user of the database.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

When you revoke CONSOLIDATE privilege from a user, SQL Anywhere:

- Removes the user from the SQL Remote system.
- Reverts that user or group to a normal user/group.
- Unsubscribes the user or group from all publications.

### Procedure

1. In SQL Central, use the [SQL Anywhere 17](#) plug-in to connect to the database.
2. Double-click the database name.
3. Double-click *Users*.
4. Right-click the consolidated user or group and click *Revoke Consolidated*.

### Results

The user/group is reverted to a normal user/group, removed from the SQL Remote system, and unsubscribed from all publications.

### Related Information

[Granting CONSOLIDATE Privilege \(SQL Central\) \[page 33\]](#)

[REVOKE CONSOLIDATE Statement \[SQL Remote\]](#)

## 1.2.2.6 Replication-related System Roles for SQL Remote

There are two replication-related SQL Anywhere system roles that control access and capabilities in SQL Remote: SYS\_RUN\_REPLICATION\_ROLE, and SYS\_REPLICATION\_ADMIN\_ROLE.

### SYS\_RUN\_REPLICATION\_ROLE System Role

Users with the SYS\_RUN\_REPLICATION\_ROLE system role have full administrative privileges on the database only when connecting from SQL Remote. SQL Remote users have full access to the database and can make changes specified in the messages.

Only users with SYS\_RUN\_REPLICATION\_ROLE system role can run SQL Remote. The SYS\_RUN\_REPLICATION\_ROLE system role has the following properties:

#### No distinct privileges when not connected from SQL Remote

A user granted the SYS\_RUN\_REPLICATION\_ROLE system role cannot exercise the privileges inherited from this role on any connection apart from SQL Remote or dbmlsync. Therefore, even if the user name and password for a user with the SYS\_RUN\_REPLICATION\_ROLE system role are widely distributed, there are no security issues. As long as the user name has no privileges beyond CONNECT granted on the database, no one can use this user name to access data in the database.

#### Full database administration privileges from SQL Remote

When connecting from SQL Remote, a user with SYS\_RUN\_REPLICATION\_ROLE system role has full database administration privileges on the database.

When you create users on the consolidated database, grant the SYS\_RUN\_REPLICATION\_ROLE system role to the publisher of the consolidated database and to each remote user. When a remote database is extracted by the Extraction utility (dbxtract) or the *Extract Database Wizard*, the remote user becomes the publisher of the remote database, and is granted the PUBLISH privilege and the SYS\_RUN\_REPLICATION\_ROLE system role.

This recommendation simplifies the administration of users. Each remote user only needs one user name to connect to the database, whether from the SQL Remote (which provides the user with the SYS\_RUN\_REPLICATION\_ROLE system role) or from any other client application (in which case the SYS\_RUN\_REPLICATION\_ROLE system role grants the user no extra privileges).

#### i Note

GRANT REMOTE DBA is a deprecated statement. Use GRANT ROLE SYS\_RUN\_REPLICATION\_ROLE instead.

### SYS\_REPLICATION\_ADMIN\_ROLE System Role

Users with the SYS\_REPLICATION\_ADMIN\_ROLE system role can perform administrative tasks related to replication, such as granting replication roles, managing publications, subscriptions, synchronization users and profiles, managing message types, and setting replication-related options.

## Related Information

[Replication-related System Roles](#)  
[Granting a Role \(SQL Central\)](#)  
[Revoking a Role \(SQL Central\)](#)  
[GRANT ROLE Statement](#)  
[REVOKE ROLE Statement](#)

### 1.2.3 Subscriptions

You subscribe a user to a publication by creating a **subscription**. Each database that shares information in a publication must have a subscription to the publication.

Periodically, the changes made to each publication in a database are replicated to all subscribers of that publication. These replications are called **publication updates**

To subscribe a user to a publication, you need the following information:

#### **Publication name**

The name of the publication to which the user is being subscribed.

#### **Subscription value**

The subscription value only applies if your publication includes a SUBSCRIBE BY clause. The subscription value is the value that is tested against the SUBSCRIBE BY clause of the publication. For example, if a publication has the name of a column containing an employee ID as a SUBSCRIBE BY clause, the value of the employee ID of the subscribing user must be provided when the subscription is created. The subscription value is always a string.

This value is only needed when the publication has a SUBSCRIBE BY clause.

#### **Subscriber-id**

The user who is being subscribed to the publication. At the consolidated database, when you create a subscription to a remote user, the remote user must have been granted REMOTE privilege. At the remote database, when you create a subscription to the consolidated user, that user must have been granted CONSOLIDATED privilege.

By default, the Extraction utility (dbxtract) and the *Extract Database Wizard* grant the appropriate PUBLISH and CONSOLIDATE privileges to users in the remote databases.

#### **In this section:**

[Creating a SQL Remote Subscription \(SQL Central\) \[page 37\]](#)  
Subscribe a user to a publication by creating a subscription.

## Related Information

[Subscription Resynchronization \[page 160\]](#)

[Publishing Only Some Rows Using the SUBSCRIBE BY Clause \(SQL Central\) \[page 18\]](#)

[Starting Subscriptions \(SQL Central\) \[page 164\]](#)

[Stopping Subscriptions \(SQL Central\) \[page 165\]](#)

[CREATE SUBSCRIPTION Statement \[SQL Remote\]](#)

[GRANT REMOTE Statement \[SQL Remote\]](#)

## 1.2.3.1 Creating a SQL Remote Subscription (SQL Central)

Subscribe a user to a publication by creating a subscription.


### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role and the publication must already have been created.

### Context

Each database that shares information in a publication must have a subscription to the publication. Periodically, the changes made to each publication in a database are replicated to all subscribers of that publication. These replications are called publication updates

### Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click the database name.
3. Double-click the *SQL Remote Subscriptions*.
4. Click **File** > **New** > *SQL Remote Subscription* .
5. Follow the instructions in the *Create SQL Remote Subscription Wizard*.

The details of the subscription are different depending on whether the publication uses a subscription expression.

### Results

The user is subscribed to the publication.

## 1.2.4 Transaction Log-based Replication

SQL Remote replicates committed changes and changes that modify data belonging to publications.

### Committed changes

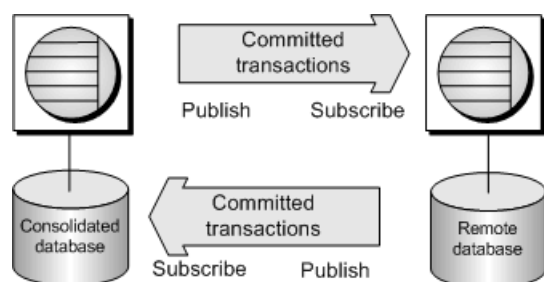
Changes that have been made to databases as recorded in their transaction log.

### Changes that modify data belonging to publications

SQL Remote scans the transaction log for committed changes to rows that belong to publications, packages the SQL statements into messages, and sends them to the subscribed databases.

On the consolidated database, all committed transactions in the transaction log that belong to publications are sent periodically to the remote databases.

On the remote databases, all committed transactions in the transaction log that belong to publications are sent periodically to the consolidated database.



## How the Database Server Handles Publications

The database server is the component that evaluates the publications and writes the information to the transaction log. The more publications you have, the more work the database server must do.

The database server evaluates the subscription expression for each update made to a table that is part of a publication. It adds the value of the expression to the transaction log, both before and after the update. For a table that is part of more than one publication, the subscription expression is evaluated before and after the update for each publication.

The additional information in the transaction log can affect performance in the following cases:

### Expensive expressions

When a subscription expression is expensive to evaluate, it can affect performance.

### Many publications

When a table belongs to several publications, many expressions must be evaluated. In contrast, the number of *subscriptions* is irrelevant to the database server.

### Many-valued expressions

Some expressions are many-valued, which can lead to additional information in the transaction log. This can affect performance.

## Subscriptions are Handled by SQL Remote

SQL Remote is the component that carries out the replication of statements.

During the send phase, the SQL Remote Message Agent maps the current subscriptions to the publication information in the transaction log and generates the appropriate messages for each remote user.

### In this section:

#### [INSERT and DELETE Statement Replication \[page 39\]](#)

INSERT and DELETE statements are treated differently on consolidated and remote databases.

#### [UPDATE Statement Replication \[page 40\]](#)

UPDATE statements are replicated differently depending on the effect they have on the affected row in the database.

#### [Procedure Replication \[page 42\]](#)

SQL Remote replicates procedures by replicating the actions of a procedure.

#### [Trigger Replication \[page 42\]](#)

Typically, remote databases have the same triggers defined as the consolidated database does.

#### [Data Definition Statements \[page 43\]](#)

Data definition statements (CREATE, ALTER, and DROP) are not replicated by SQL Remote unless they are executed in passthrough mode.

#### [SQL Remote Data Types \[page 44\]](#)

SQL Remote does not perform any character set conversions.

## Related Information

[SQL Remote Performance \[page 102\]](#)

[The Transaction Log](#)

### 1.2.4.1 INSERT and DELETE Statement Replication

INSERT and DELETE statements are treated differently on consolidated and remote databases.

- When an INSERT statement is executed on one database, it is sent as an INSERT statement to the subscribed databases in the SQL Remote system.
- When a DELETE statement is executed on one database, it is sent as a DELETE statement to the subscribed databases in the SQL Remote system.

## Consolidated Database

SQL Remote copies each INSERT or DELETE statement from the consolidated database transaction log, and sends them to each remote database that subscribes to the row being inserted or deleted. When only a subset

of the columns in the table is subscribed to, the INSERT statements sent to the remote databases contain only those columns.

## Remote Databases

SQL Remote copies each INSERT or DELETE statement from a remote database transaction log and sends it to the consolidated database that subscribes to the row being inserted or deleted. The consolidated database then applies the statement, which results in writing to its transaction log. When the consolidated database transaction log is processed by SQL Remote, the changes are eventually propagated to the other remote sites. SQL Remote ensures that statements are not sent to the remote user that initially executed them.

## Related Information

[Duplicate Primary Key Errors \[page 58\]](#)

[Row Not Found Errors \[page 55\]](#)

### 1.2.4.2 UPDATE Statement Replication

UPDATE statements are replicated differently depending on the effect they have on the affected row in the database.

The following scenarios describe how an UPDATE statement is replicated

- When an UPDATE statement has the effect of updating a row in a given remote user's subscription, it is sent to that user as an UPDATE statement.
- When an UPDATE statement has the effect of removing a row from a given remote user's subscription, it is sent to that user as a DELETE statement.
- When an UPDATE statement has the effect of adding a row to a given remote user's subscription, it is sent to that user as an INSERT statement.

To demonstrate how an UPDATE statement can be replicated, the following example uses a consolidated database and three remote databases for the users: Ann, Marc, and ManagerSteve.

Consolidated			Ann	Marc	ManagerSteve	
ID	Rep	Dept	ID	Rep	ID	Rep
1	Ann	101	1	Ann	1	Ann
2	Marc	101		2	2	Marc
3	Marc	101		3	3	Marc
4	Rob	102				



On the consolidated database, there is a publication, named cons, that is created with the following statement:

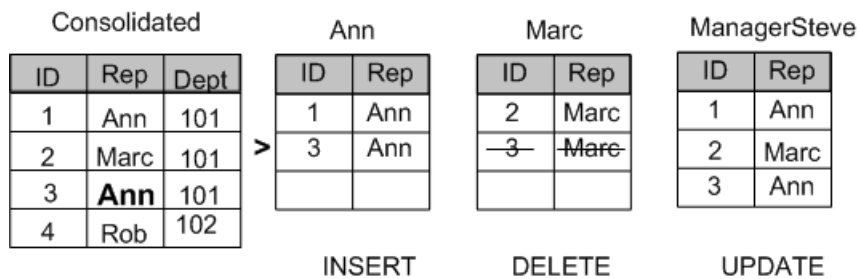
```
CREATE PUBLICATION "cons"."p1" (
  TABLE "DBA"."customers" ( "ID", "Rep") SUBSCRIBE BY repid
);
```

Ann and Marc subscribe to the cons publication with their respective Rep column values. ManagerSteve subscribes to the cons publication with both Ann and Marc's Rep column values. The following statements subscribe the three users to the publication cons:

```
CREATE SUBSCRIPTION
  TO "cons"."p1" ( 'Ann' )
  FOR "Ann";
CREATE SUBSCRIPTION
  TO "cons"."p1" ( 'Marc' )
  FOR "Marc";
CREATE SUBSCRIPTION
  TO "cons"."p1" ( 'Ann' )
  FOR "ManagerSteve";
CREATE SUBSCRIPTION
  TO "cons"."p1" ( 'Marc' )
  FOR "ManagerSteve";
```

On the consolidated database, an UPDATE statement that changes the Rep value of a row from Marc to Ann is replicated to:

- Marc as a DELETE statement.
- Ann as an INSERT statement.
- ManagerSteve as an UPDATE statement.



This reassignment of rows among subscribers is sometimes called **territory realignment** because it is a common feature of sales force automation applications, where customers are periodically reassigned among representatives.

## Related Information

- [Update Conflicts \[page 47\]](#)
- [UPDATE Statement](#)

## 1.2.4.3 Procedure Replication

SQL Remote replicates procedures by replicating the actions of a procedure.

The procedure call is not replicated. Instead, the individual actions (INSERT, UPDATE, and DELETE statements) of the procedure are replicated.

## 1.2.4.4 Trigger Replication

Typically, remote databases have the same triggers defined as the consolidated database does.

By default, SQL Remote does not replicate the actions performed by the triggers. Instead, when an action that fires a trigger on the consolidated database is replicated on the remote database, the duplicate trigger is automatically fired on the remote database. This avoids privileges issues and the possibility of each action occurring twice. There are some exceptions to this rule:

### Replication of RESOLVE UPDATE triggers

The actions carried out by conflict resolution, or RESOLVE UPDATE triggers *are* replicated from the consolidated database to all remote databases, including the remote database that sent the message that created the conflict.

### Replication of BEFORE triggers

The actions of a BEFORE trigger that modifies a row being updated *are* replicated before the UPDATE statement actions. For example, a BEFORE UPDATE trigger that increases a counter column in the row to keep track of the number of times a row is updated would double count if replicated as the BEFORE UPDATE trigger fires on the remote database when the UPDATE statement is replicated.

A BEFORE UPDATE trigger that sets a column to the time of the last update also receives the time the UPDATE statement is replicated.

To prevent this problem, you must ensure that, at the subscriber database, the BEFORE UPDATE trigger is not present or does not perform the replicated action.

## An Option to Replicate Trigger Actions

To replicate all trigger actions when sending messages, use the -t option for the SQL Remote Message Agent (dbremote).

When you use the -t option, ensure that the trigger actions are not carried out twice at remote databases, once when the replicated trigger actions are applied, and once when the trigger is fired on the remote database.

To ensure that trigger actions are not carried out twice, use one of the following options:

- Wrap an IF CURRENT REMOTE USER IS NULL ... END IF statement around the body of the trigger.
- Set the SQL Anywhere fire\_triggers option to Off for the SQL Remote user name.

## Avoiding Trigger Errors

If a publication includes only a subset of a database, a trigger at the consolidated database can refer to tables or rows that are present at the consolidated database, but are not present on the remote databases. When such a trigger is fired on the remote database, errors occur. To avoid these errors, use an IF statement to make the trigger actions conditional, and:

- Have the actions of the trigger be conditional on the value of CURRENT PUBLISHER.
- Have the actions of the trigger be conditional on the object\_id function not returning NULL. The object\_id function takes a table or other object as an argument, and returns the ID number of that object or NULL if the object does not exist.
- Have actions of the trigger be conditional on a SELECT statement that determines if the rows exist.

## Extraction Utility (dbxtract)

By default, the database Extraction utility (dbxtract) and the *Extract Database Wizard* extract the trigger definitions.

## Related Information

[The CURRENT REMOTE USER Special Value \[page 53\]](#)

[System Functions](#)

[Default Resolution for Update Conflicts \[page 48\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

[fire\\_triggers Option](#)

[CURRENT PUBLISHER Special Value](#)

## 1.2.4.5 Data Definition Statements

Data definition statements (CREATE, ALTER, and DROP) are not replicated by SQL Remote unless they are executed in passthrough mode.

## Related Information

[SQL Remote Passthrough Mode \[page 157\]](#)

## 1.2.4.6 SQL Remote Data Types

SQL Remote does not perform any character set conversions.

### Use Compatible Sort Orders and Character Sets

The character set and collation used by the SQL Anywhere consolidated database must be the same as the remote database's.

#### In this section:

[BLOBs \[page 44\]](#)

**BLOBs** include the LONG VARCHAR, LONG BINARY, TEXT, and IMAGE data types.

[SQL Remote Dates and Times \[page 46\]](#)

When date or time columns are replicated, SQL Remote uses the setting of the `sr_date_format`, `sr_time_format`, and `sr_timestamp_format` database options to format the date.

### Related Information

[International Languages and Character Sets](#)

## 1.2.4.6.1 BLOBs

**BLOBs** include the LONG VARCHAR, LONG BINARY, TEXT, and IMAGE data types.

When SQL Remote replicates an INSERT or UPDATE statement, it uses a variable in place of the BLOB value. That is, the BLOB is broken into pieces and replicated in chunks. At the recipient database, the pieces are reconstituted by using a SQL variable and concatenated. The value of the variable is built up by a sequence of statements of the form:

```
SET var = var || 'more_stuff';
```

The variable makes the size of the SQL statements involving long values smaller, so they fit within a single message.

The SET statements are separate SQL statements, so that the BLOB is effectively split over several SQL Remote messages.

## Controlling Replication of BLOBs

The SQL Anywhere `blob_threshold` option allows further control over the replication of long values. Any value that is longer than the `blob_threshold` option is replicated as though it is a BLOB value.

## Using the `verify_threshold` Option to Minimize Message Size

The `verify_threshold` database option can prevent long values from being verified (in the `VERIFY` clause of a replicated `UPDATE` statement). The default value for the option is 1000. When the data type of a column is longer than the threshold, old values for the column are not verified when an `UPDATE` statement is replicated. This reduces the size of SQL Remote messages, but has the disadvantage that conflicting updates of long values are not detected.

Use the following technique to detect conflicts when the `verify_threshold` option is being used to reduce the size of messages:

1. Configure your databases so that whenever a BLOB is updated, a `last_modified` column in the same table is also updated.
2. Configure your publications so that the `last_modified` column is replicated with the BLOB column.
3. When the BLOB column and the `last_modified` column are replicated, the values in the `last_modified` column can be verified. If there is a conflict with the `last_modified` column, then there is a conflict with the BLOB column as well.

## Using a Work Table to Avoid Redundant Updates

Repeated updates to a BLOB should be done in a work table, and the final version should be assigned to the replicated table. For example, if a document in progress is updated 20 times throughout the day and SQL Remote is run once at the end of the day, all 20 updates are replicated. If the document is 200 KB in length, then 4 MB of messages are sent.

Use a `document_in_progress` table. When the user is done revising a document, the application moves it from the `document_in_progress` table to the replicated table. This results in a single update (200 KB of messages).

## Related Information

[SET Statement](#)

[blob\\_threshold Option \[SQL Remote\]](#)

[verify\\_threshold Option \[SQL Remote\]](#)

## 1.2.4.6.2 SQL Remote Dates and Times

When date or time columns are replicated, SQL Remote uses the setting of the `sr_date_format`, `sr_time_format`, and `sr_timestamp_format` database options to format the date.

For example, the following option setting instructs SQL Remote to send a date of May 2, 1998 as 1998-05-02.

```
SET OPTION sr_date_format = 'yyyy-mm-dd';
```

When replicating dates and times:

- The time, date, and timestamp formats must be consistent throughout the SQL Remote system.
- The order of the year, month, and day used for the date and timestamp formats must match the setting of the `date_order` database option.
- The `date_order` option for the duration of each connection can be changed.

### Related Information

[sr\\_date\\_format Option \[SQL Remote\]](#)

[sr\\_time\\_format Option \[SQL Remote\]](#)

[sr\\_timestamp\\_format Option \[SQL Remote\]](#)

[date\\_order Option](#)

## 1.2.5 Replication Conflicts and Errors

SQL Remote allows databases to be updated at multiple databases. Careful design is required to avoid replication errors, especially when the database has a complicated structure.

### Replication Conflicts

Replication conflicts are different from replication errors. SQL Remote provides appropriate conflict resolution as part of the regular operation of a SQL Remote system, using triggers and procedures.

### Replication Errors

Replication errors fall into the following categories:

#### Row not found errors

A user deletes a row (with a given primary key value.) A second user updates or deletes the same row at another site. In this case, the second statement fails, as the row is not found.

#### Referential integrity errors

When a column containing a foreign key is included in a publication, but the associated primary key is not included, INSERT statements that reference the foreign key fail.

Also, referential integrity errors can occur when a primary table has a SUBSCRIBE BY expression and the associated foreign table does not: rows from the foreign table may be replicated, but the rows from the primary table may be excluded from the publication.

### **Duplicate primary key errors**

Two users insert a row using the same primary key values, or one user updates a primary key and a second user inserts a primary key of the new value. The second operation to reach a given database in the replication system fails because it would produce a duplicate primary key.

## **Delivery Errors**

Delivery errors are managed using the SYSREMOTEUSER system table, part of the Guaranteed Message Delivery System.

## **Related Information**

[Replication Error-handling Procedures \[page 150\]](#)

[Guaranteed Message Delivery System \[page 114\]](#)

[Custom Conflict Resolution Using Triggers \[page 52\]](#)

[Row Not Found Errors \[page 55\]](#)

[Referential Integrity Errors \[page 56\]](#)

[Duplicate Primary Key Errors \[page 58\]](#)

[Default Resolution for Update Conflicts \[page 48\]](#)

[Replication Error Reporting and Handling \[page 149\]](#)

[Update Conflicts \[page 47\]](#)

## **1.2.6 Update Conflicts**

Update conflicts cannot occur when data is shared for reading, or when each row (as identified by its primary key) is updated at only one database. Update conflicts only occur when data is updated at more than one database.

To replicate UPDATE statements, SQL Remote issues a separate UPDATE statement for each row. These single-row statements can fail for one of the following reasons:

### **The row to be updated differs in one or more of its columns**

When one of the values expected to be present has been changed by some other user, an **update conflict** occurs.

On remote databases, the update takes place regardless of the values in the row.

On the consolidated database, SQL Remote allows **conflict resolution** operations to take place. For example, when a conflict is detected, the consolidated database can:

- Use the default conflict resolution.
- Use a customized conflict resolution that uses the VERIFY clause.
- Use a customized conflict resolution that uses triggers.

#### **i** Note

UPDATE statement conflicts do *not* apply to primary key updates. You should not update primary keys in a SQL Remote system. Primary key conflicts must be excluded from the system using proper design.

#### **The row to be updated does not exist**

Each row is identified by its primary key values. If the row has been deleted or if a primary key has been altered by another user, the row to be updated cannot be found.

On remote databases, the update does not occur.

On the consolidated database, the update does not occur.

#### **A table without a primary key or unique constraint refers to all columns in the WHERE clause of replicated updates**

When two remote databases make separate updates to the same row and replicate the changes to the consolidated database, the first changes to arrive on the consolidated database are applied; changes from the second database are not applied. As a result, databases become inconsistent. All replicated tables should have a primary key or a unique constraint and the columns in the constraint should never be updated.

#### **In this section:**

##### [Default Resolution for Update Conflicts \[page 48\]](#)

Update conflicts only occur when data is being updated at more than one site.

##### [Custom Conflict Resolution Using a VERIFY Clause \[page 49\]](#)

SQL Remote generates UPDATE statements in the messages that use the VERIFY clause. An UPDATE statement changes the value of one or more rows from the existing value to a new value.

##### [Custom Conflict Resolution Using Triggers \[page 52\]](#)

Custom conflict resolution can take several forms. For example, in some applications, resolution can:

## **1.2.6.1 Default Resolution for Update Conflicts**

Update conflicts only occur when data is being updated at more than one site.

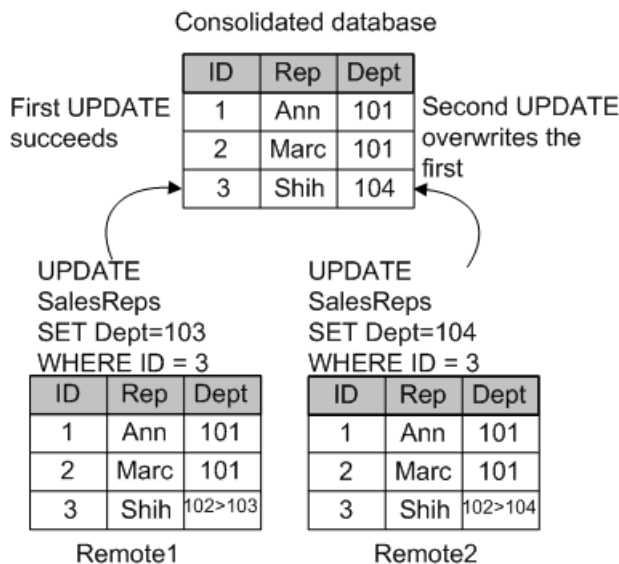
For example:

1. User 1 updates a row at remote site 1.
2. User 2 updates the same row at remote site 2.
3. The update from User 1 is sent and applied to the consolidated database.
4. The update from User 2 is sent to the consolidated database.



The **default** method for resolving this type of update conflict is the following:

1. The more recent operation (from User 2) succeeds. It becomes the value in the consolidated database and it is the value that is replicated to all other databases subscribed to that row.
2. All other updates (from User 1) are lost.
3. No report is made of the conflict.



## Related Information

[Row Not Found Errors \[page 55\]](#)

### 1.2.6.2 Custom Conflict Resolution Using a VERIFY Clause

SQL Remote generates UPDATE statements in the messages that use the VERIFY clause. An UPDATE statement changes the value of one or more rows from the existing value to a new value.

UPDATE statements that include a VERIFY clause also contain the existing value of the row.

When applying an UPDATE statement, the consolidated database compares its existing value of the row with what the remote database expects the existing value of the row to be. An update conflict is detected by the database server when the VERIFY clause values don't match the rows in the database.

For example, an update conflict occurs when the following sequence of events takes place:

1. User 1 updates a row at remote site 1.
2. User 2 updates the same row at remote site 2.
3. The update from User 1 is sent and applied to the consolidated database.
4. The update from User 2 is sent to the consolidated database.

Because the UPDATE statement contains a VERIFY clause, the consolidated database can detect conflicts. On the consolidated database, SQL remote compares the value in its row with the old row value that User 2 sent. As these values are not the same, there is an update conflict.

When an update conflict is detected, the consolidated database:

1. Fires any conflict resolution triggers defined for the operation.  
You define **conflict resolution triggers** to handle update conflicts. Conflict resolution triggers are fired only at a consolidated database when messages are applied by a remote user.
2. Executes the UPDATE statements.
3. Sends any actions of the conflict resolution trigger and the UPDATE statement to all remote databases, including the sender of the message that triggered the conflict.  
Typically, SQL Remote does not replicate the actions of triggers; the trigger is assumed present on the remote database. Conflict resolution triggers are fired only on consolidated databases, and so their actions are replicated to remote databases.
5. The remote databases receives the UPDATE statements from the consolidated database.  
On remote databases, RESOLVE UPDATE triggers are not fired when a message from a consolidated database contains an update conflict.
6. On the remote database, the UPDATE statements are processed.  
At the end of the process, the data is consistent throughout the system.

#### In this section:

##### [UPDATE Statements with a VERIFY Clause \[page 50\]](#)

SQL Remote generates UPDATE statements in the messages that use the VERIFY clause. An UPDATE statement changes the value of one or more rows from the existing value to a new value.

##### [The verify\\_all\\_columns Option \[page 51\]](#)

By default, the database option verify\_all\_columns is Off. When it is set to Off, only those columns that are updated are checked.

## Related Information

### [Custom Conflict Resolution Using Triggers \[page 52\]](#)

## 1.2.6.2.1 UPDATE Statements with a VERIFY Clause

SQL Remote generates UPDATE statements in the messages that use the VERIFY clause. An UPDATE statement changes the value of one or more rows from the existing value to a new value.

An UPDATE statement with a VERIFY clause takes the following form:

```
UPDATE table-list
SET column-name = expression, ...
[ VERIFY (column-name, ...)
  VALUES ( expression, ...) ]
[ WHERE search-condition ]
[ ORDER BY expression [ ASC | DESC ], ... ]
```

The VERIFY clause can be used only when the `table-list` parameter consists of a single table. It compares the values of specified columns to a set of expected values, which are the values that were present in the publisher database when the UPDATE statement was applied. When the VERIFY clause is specified, only one table can be updated at a time.

The VERIFY clause is useful only for single-row updates. However, a multi-row UPDATE statement executed on a database is replicated as a set of single-row UPDATE statements by the SQL Remote, so this imposes no constraints on client applications.

## Related Information

[UPDATE Statement \[SQL Remote\]](#)

### 1.2.6.2.2 The `verify_all_columns` Option

By default, the database option `verify_all_columns` is Off. When it is set to Off, only those columns that are updated are checked.

When the `verify_all_columns` option is set to On, then:

- All columns are verified on replicated UPDATE statements.
- A RESOLVE UPDATE trigger is fired whenever any column is different.
- The size of messages is bigger, because more information is sent for each UPDATE statement.

You can set the `verify_all_columns` option either for the PUBLIC role or just for the user contained in the SQL Remote connection string.

## The Extraction Utility (`dbxtract`)

When the `verify_all_columns` option is set on the consolidated database before the remote databases are extracted, then the `verify_all_columns` option on the remote databases is set by the Extraction utility (`dbxtract`) and the [Extract Database Wizard](#).

## Related Information

[verify\\_all\\_columns Option \[SQL Remote\]](#)

## 1.2.6.3 Custom Conflict Resolution Using Triggers

Custom conflict resolution can take several forms. For example, in some applications, resolution can:

- Compare the dates of the original transactions.
- Perform calculations on the results of two or more updates.
- Report the conflict into a table.

Custom conflict resolution requires you to write RESOLVE UPDATE triggers.

### Using the RESOLVE UPDATE Conflict Resolution Trigger

RESOLVE UPDATE triggers fire before each row is updated. The syntax for a RESOLVE UPDATE trigger is as follows:

```
CREATE TRIGGER trigger-name
RESOLVE UPDATE
OF column-name ON table-name
[ REFERENCING [ OLD AS old-val ]
  [ NEW AS new-val ]
  [ REMOTE AS remote-val ] ]
FOR EACH ROW
BEGIN
  ...
END
```

The REFERENCING clause allows access to the values in the row of the table to be updated (OLD), to the values the row is to be updated to (NEW), and to the rows that should be present according to the VERIFY clause (REMOTE). Only columns present in the VERIFY clause can be referenced in the REMOTE AS clause; other columns return an error.

#### In this section:

##### [The CURRENT REMOTE USER Special Value \[page 53\]](#)

The CURRENT REMOTE USER special value is set by the receive phase of SQL Remote when it is applying messages to the database.

##### [Date Conflict Resolution \[page 53\]](#)

To illustrate how you can resolve date conflicts, suppose a table in a contact management system has a column holding the more recent contact with each customer.

##### [Inventory Conflict Resolution \[page 54\]](#)

Consider a warehouse system for a manufacturer of sporting goods. There is a table of product information, with a Quantity column holding the remaining quantity of each product. An update to this column typically depletes the quantity in stock or, if a new shipment is brought in, adds to it.

## Related Information

[CREATE TRIGGER Statement](#)

[UPDATE Statement \[SQL Remote\]](#)

[Trigger Replication \[page 42\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

### 1.2.6.3.1 The CURRENT REMOTE USER Special Value

The CURRENT REMOTE USER special value is set by the receive phase of SQL Remote when it is applying messages to the database.

The CURRENT REMOTE USER special value is most useful in triggers to determine whether the operations being applied are being applied by the receive phase of SQL Remote, and if they are, which remote user generated the operations being applied.

### 1.2.6.3.2 Date Conflict Resolution

To illustrate how you can resolve date conflicts, suppose a table in a contact management system has a column holding the more recent contact with each customer.

One representative talks with a customer on Friday, but does not upload his changes to the consolidated database until Monday. Meanwhile, a second representative meets the customer on Saturday, and updates the changes that evening.

There is no conflict when the Saturday update is replicated to the consolidated database, but when the Monday update arrives, it finds the row already changed.

By default, the Monday update would proceed, leaving the column with the incorrect information that the more recent contact occurred on Friday. However, update conflicts on this column should be resolved by inserting the more recent date in the row.

## Implementing the Solution

The following RESOLVE UPDATE trigger chooses the more recent of the two new values and enters it in the database.


```
CREATE TRIGGER contact_date RESOLVE UPDATE
ON Contacts
REFERENCING OLD AS old_name
NEW AS new_name
FOR EACH ROW
BEGIN
  IF new_name.contact_date <
    old_name.contact_date THEN
    SET new_name.contact_date
      = old_name.contact_date
  END IF
END;
```

If the value being updated is later than the value that would replace it, the new value is reset to leave the entry unchanged.

### 1.2.6.3.3 Inventory Conflict Resolution


Consider a warehouse system for a manufacturer of sporting goods. There is a table of product information, with a Quantity column holding the remaining quantity of each product. An update to this column typically depletes the quantity in stock or, if a new shipment is brought in, adds to it.

A sales representative at a remote database enters an order, depleting the stock of small, tank top, tee shirts by five, from 28 to 23, and enters this in her database. Meanwhile, before this update is replicated to the consolidated database, another sales representative receives 40 returned tee shirts. This sales representative enters the returns into his remote database and replicates the changes to the consolidated database at the warehouse, adding 40 to the Quantity column to make it 68.




The warehouse entry is added to the database: the Quantity column now shows there are 68 small tank-top tee shirts in stock. When the update from the first sales representative arrives, it causes a conflict: SQL Anywhere detects that the update is from 28 to 23, but that the current value of the column is 68.

By default, the more recent update succeeds, and the inventory level is set to the incorrect value of 23.



In this example, the conflict should be resolved by summing the changes to the inventory column to produce the result, so that a final value of 63 is placed into the database.



## Implementing the Solution

A suitable RESOLVE UPDATE trigger for this situation would add the increments from the two updates. For example:

```
CREATE TRIGGER resolve_quantity
RESOLVE UPDATE OF Quantity
ON "DBA".Products
REFERENCING OLD AS old_name
NEW AS new_name
REMOTE AS remote_name
FOR EACH ROW
BEGIN
    SET new_name.Quantity = new_name.Quantity
                        + old_name.Quantity
                        - remote_name.Quantity
END;
```

This trigger adds the difference between the old value in the consolidated database (68) and the old value in the remote database when the original UPDATE statement was executed (28) to the new value being sent, before the UPDATE statement is implemented. So, new\_name.Quantity becomes 63 (= 23 + 68 - 28), and this value is entered into the Quantity column.

Consistency is maintained on the remote database as follows:

1. The original remote UPDATE statement changed the value from 28 to 23.
2. The warehouse's entry is replicated to the remote database, but fails as the old value is not what was expected.
3. The changes made by the RESOLVE UPDATE trigger are replicated to the remote database.

## 1.2.7 Row Not Found Errors

A user deletes a row (with a given primary key value.) A second user updates or deletes the same row at another site. In this case, the second statement fails, as the row is not found.

To replicate UPDATE and DELETE statements correctly, you must include all of the primary key columns in the article.

When an UPDATE or a DELETE statement is replicated, SQL Remote uses the primary key columns to uniquely identify the row being updated or deleted. All tables being replicated must have a declared primary key or unique constraint. A unique index is not enough.

## WHERE Clause and Primary Keys

The primary key columns are used in the WHERE clause of replicated UPDATE and DELETE statements. When a table has no primary key, the WHERE clause refers to all columns in the table.

## Related Information

[Replication Error Reporting and Handling \[page 149\]](#)

[Default Resolution for Update Conflicts \[page 48\]](#)

[INSERT and DELETE Statement Replication \[page 39\]](#)

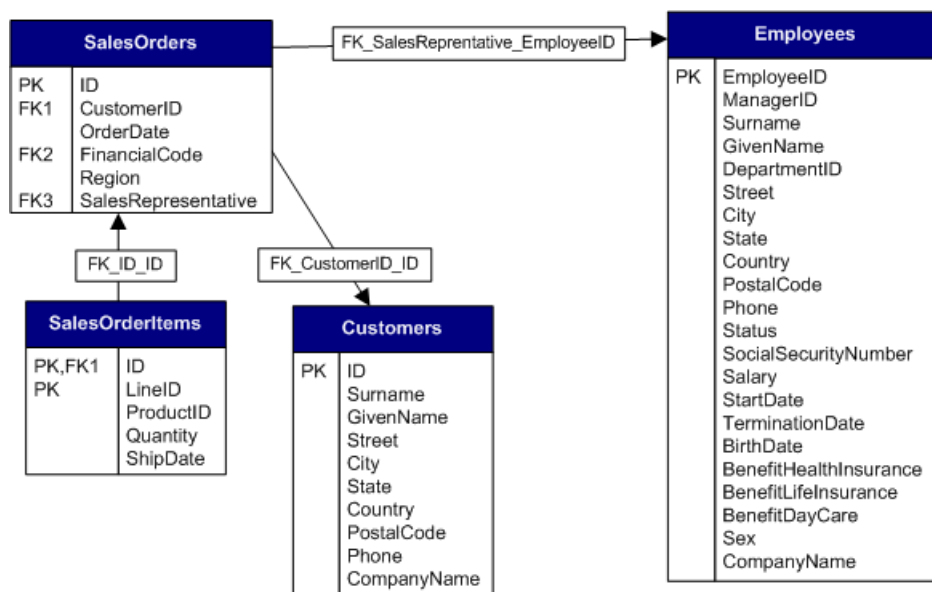
## 1.2.8 Referential Integrity Errors

The tables in a relational database are often related through foreign key references. As a result, referential integrity constraints ensure that the database remains consistent.

When you replicate only a part of a database, you must ensure that the replicated database still has referential integrity.

You want to avoid unreplicated referenced table errors. Your remote databases should not contain foreign keys that point to unreplicated tables.

For example, in a consolidated database the SalesOrders table has a foreign key to the Employees table. SalesOrders.SalesRepresentative is the foreign key that references the primary key, Employees.EmployeeID.





A publication, PubSales, is created that excludes the Employees table, but includes the entire SalesOrder table.

```
CREATE PUBLICATION PubSales (  
    TABLE Customers,  
    TABLE SalesOrders,  
    TABLE SalesOrderItems,  
);
```

A remote user, Rep1, subscribes to the PubSales publication. Then, you extract Rep1 from the consolidated database and try to create a database for Rep1. However, the database creation fails because Rep1 is missing the Employees table. To avoid this problem, you can:

#### **Remove the foreign key reference**

To exclude foreign key references, specify the -xf option when using the Extraction utility (dbxtract).

However, if you remove the foreign key reference from the remote database, then there is no constraint in the remote database to prevent an invalid value from being inserted into the SalesRepresentative column of the SalesOrders table.

If an invalid value is inserted in the SalesRepresentative column at the remote database, the replicated INSERT statement fails on the consolidated database.

#### **Include the missing table in the publication**

Include the Employees table (or at least its primary key) in the publication. For example:

```
CREATE PUBLICATION PubSales (  
    TABLE Customers,  
    TABLE SalesOrders,  
    TABLE SalesOrderItems,  
    TABLE Products,  
    TABLE Employees  
);
```

#### **In this section:**

##### [Errors on Insert \[page 58\]](#)

When replicating INSERT statements from a remote database to a consolidated database, you can only exclude some columns from the publication or else the inserts fail.

## **Related Information**

[Entity and Referential Integrity](#)

[Replication Error Reporting and Handling \[page 149\]](#)

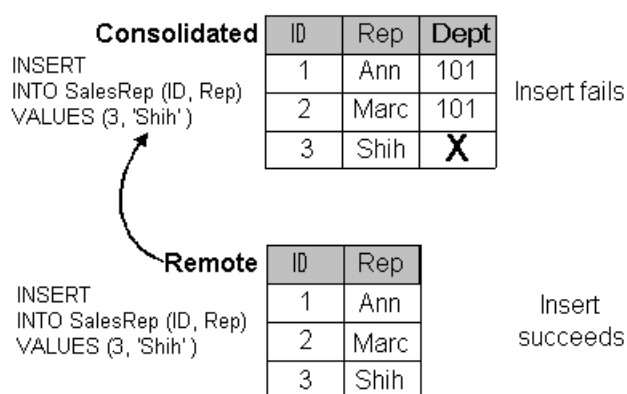
## 1.2.8.1 Errors on Insert

When replicating INSERT statements from a remote database to a consolidated database, you can only exclude some columns from the publication or else the inserts fail.

You can exclude:

- Columns that allow NULL.
- Columns that have defaults.

If you exclude any column that does not satisfy one of these requirements, then INSERT statements carried out at a remote database fail when they are replicated on the consolidated database.



### i Note

An exception to this example is when a BEFORE trigger is used to maintain the columns that are not included in the INSERT statement.

## Related Information

[Replication Conflicts and Errors \[page 46\]](#)

## 1.2.9 Duplicate Primary Key Errors

When all users are connected to the same database, there is no problem ensuring that each INSERT statement uses a unique primary key. If a user tries to re-use a primary key, the INSERT statement fails.

The situation is different in a replication system because users are connected to many databases. A potential problem arises when two users, connected to different remote databases, insert a row using the same primary key value. Each of their statements succeeds because the primary key value is unique on each remote database.

However, when these two users replicate their databases with the same consolidated database, a problem arises. The first database to replicate with the consolidated database succeeds. However, the second insert to reach a given database in the replication system fails.

## Primary Key Values Must Be Unique

To avoid primary key errors, you must ensure that when a database inserts a row, its primary key is guaranteed to be unique across all databases in the system. There are several techniques for achieving this goal, including:

1. Using the default GLOBAL AUTOINCREMENT feature of SQL Anywhere databases.
2. Using a primary key pool to maintain a list of unused, unique primary key values at each site.

These techniques can be used either separately or together to avoid duplicate primary keys.

### In this section:

[GLOBAL AUTOINCREMENT Columns in SQL Remote \[page 59\]](#)

Use the GLOBAL AUTOINCREMENT default to assign each remote database a unique global database identification number.

[Primary Key Pools with SQL Remote \[page 61\]](#)

A **primary key pool** is a table that contains a set of primary key values for each database in a SQL Remote system.

## Related Information

[Replication Error Reporting and Handling \[page 149\]](#)

### 1.2.9.1 GLOBAL AUTOINCREMENT Columns in SQL Remote

Use the GLOBAL AUTOINCREMENT default to assign each remote database a unique global database identification number.

When you specify the GLOBAL AUTOINCREMENT default for a column, 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.

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

### In this section:

[DEFAULT GLOBAL AUTOINCREMENT Declaration \[page 60\]](#)

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

## Related Information

[The GLOBAL AUTOINCREMENT Default](#)

### 1.2.9.1.1 DEFAULT GLOBAL AUTOINCREMENT Declaration

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

## Partition Size

Optionally, the partition size can be specified in parentheses immediately following the AUTOINCREMENT keyword. The partition size may be any non-negative integer, although the partition size is generally chosen so that the supply of numbers within any one partition is rarely 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 your column is not of type INT or BIGINT, specify the partition size explicitly.

## Example

The following statement creates a table with two columns: an integer that holds a customer identification number and a character string that holds the customer's name. The identification number column, ID, uses the GLOBAL AUTOINCREMENT default and has a partition size of 5000.

```
CREATE TABLE Customers (  
  ID    INT    DEFAULT GLOBAL AUTOINCREMENT (5000),  
  name  VARCHAR(128) NOT NULL,  
  PRIMARY KEY (ID)  
);
```

## Related Information

[CREATE TABLE Statement](#)

[ALTER TABLE Statement](#)

## 1.2.9.2 Primary Key Pools with SQL Remote

A **primary key pool** is a table that contains a set of primary key values for each database in a SQL Remote system.

A master primary key pool table is created and stored on the consolidated database. Remote users subscribe to the consolidated database primary key pool table to receive their own set of primary key values. When a remote user inserts a new row into a table, they use a stored procedure to select a valid primary key from their pool. The pool is maintained by periodically running a procedure on the consolidated database that replenishes the supply.

The primary key pool technique requires the following components:

### **Primary key pool table**

On the consolidated database, you need a table to hold valid primary key values for each database in the system.

### **Replenishment procedure**

On the consolidated database, you need a stored procedure to keep the key pool table filled.

### **Sharing of key pools**

Each remote database in the system must subscribe to its own set of valid values from the consolidated database key pool table.

### **Data entry procedures**

On the remote databases, new rows are entered using a stored procedure that picks the next valid primary key value from the pool and then deletes that value from the key pool.

### **In this section:**

#### [Creating a Primary Key Pool Table \(SQL\) \[page 62\]](#)

Create a primary key pool table.

#### [Replicating the Primary Key Pool \(SQL\) \[page 63\]](#)

Create a separate publication for the primary key pool and subscribe users to it.

#### [Filling and Replenishing the Key Pool \(SQL\) \[page 64\]](#)

Replenish the contents of the primary key pool table on the consolidated database and then replicate the new primary keys to the remote databases.

#### [Using the Primary Keys from the Key Pool \(SQL\) \[page 65\]](#)

Use primary keys to form the remote user's pool of available primary keys when remote users add new customers.

## 1.2.9.2.1 Creating a Primary Key Pool Table (SQL)

Create a primary key pool table.

### Procedure

1. On the consolidated database, execute the following statement to create a primary key pool table:

```
CREATE TABLE KeyPool (  
  table_name VARCHAR(128) NOT NULL,  
  value INTEGER NOT NULL,  
  location CHAR(12) NOT NULL,  
  PRIMARY KEY (table_name, value),  
);
```

Column	Description
table_name	Holds the names of tables for which primary key pools must be maintained. For example, if new sales representatives are added only on the consolidated database, only the Customers table needs a primary key pool and this column is redundant.
value	Holds a list of primary key values. Each value is unique for each table listed in table_name.
location	An identifier for the recipient. In some systems, this can be the same as the rep_key value of the SalesReps table. In other systems, there are users other than sales representatives; in such systems, the two identifiers should be distinct.

2. To increase performance, execute the following statement to create an index on the primary key table:

```
CREATE INDEX KeyPoolLocation  
ON KeyPool (table_name, location, value);
```

### Results

The primary key pool is created.

### Related Information

[CREATE TABLE Statement](#)  
[CREATE INDEX Statement](#)

## 1.2.9.2.2 Replicating the Primary Key Pool (SQL)

Create a separate publication for the primary key pool and subscribe users to it.

### Prerequisites

The primary key pool table must already have been created.

### Context

You can either incorporate the primary key pool into an existing publication or share it as a separate publication.

### Procedure

1. On the consolidated database, create a publication for the primary key pool data.

```
CREATE PUBLICATION KeyPoolData (  
    TABLE KeyPool SUBSCRIBE BY location  
);
```

2. Create subscriptions for each remote database to the KeyPoolData publication.

```
CREATE SUBSCRIPTION  
    TO KeyPoolData( 'Sam_Singer' )  
    FOR user1;  
CREATE SUBSCRIPTION  
    TO KeyPoolData( 'user2' )  
    FOR user2;  
...
```

The subscription argument is the location identifier.

### Results

The primary key pool is replicated.

### Related Information

[CREATE PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

## 1.2.9.2.3 Filling and Replenishing the Key Pool (SQL)

Replenish the contents of the primary key pool table on the consolidated database and then replicate the new primary keys to the remote databases.

### Prerequisites

The primary key pool table must already have been created.

### Procedure

1. On the consolidated database, create a procedure to fill the primary key pool.

#### Note

You cannot use a trigger to replenish the key pool, as trigger actions are not replicated.

For example:

```
CREATE PROCEDURE ReplenishPool()
BEGIN
  FOR EachTable AS TableCursor
  CURSOR FOR
    SELECT table_name
    AS CurrTable, max(value) as MaxValue
    FROM KeyPool
    GROUP BY table_name
  DO
    FOR EachRep AS RepCursor
    CURSOR FOR
      SELECT location
      AS CurrRep, COUNT(*) AS NumValues
      FROM KeyPool
      WHERE table_name = CurrTable
      GROUP BY location
    DO
      // make sure there are 100 values.
      // Fit the top-up value to your
      // requirements
      WHILE NumValues < 100 LOOP
        SET MaxValue = MaxValue + 1;
        SET NumValues = NumValues + 1;
        INSERT INTO KeyPool
        (table_name, location, value)
        VALUES
        (CurrTable, CurrRep, MaxValue);
      END LOOP;
    END FOR;
  END FOR;
END FOR;
```



```
END;
```

2. Insert an initial primary key value in the primary key pool for each user.

The ReplenishPool procedure requires at least one primary key value to exist for each subscriber, so that it can find the maximum value and add one to generate the next set.

To initially fill the pool, you can insert a single value for each user, and then call ReplenishPool to fill up the rest. The following example illustrates this for three remote users and a single consolidated user named Office:

```
INSERT INTO KeyPool VALUES( 'Customers', 40, 'user1' );
INSERT INTO KeyPool VALUES( 'Customers', 41, 'user2' );
INSERT INTO KeyPool VALUES( 'Customers', 42, 'user3' );
INSERT INTO KeyPool VALUES( 'Customers', 43, 'Office');
CALL ReplenishPool();
```

The ReplenishPool procedure fills the pool for each user up to 100 values. The value you need depends on how often users are inserting rows into the tables in the database.

3. Periodically run ReplenishPool.

The ReplenishPool procedure must be run periodically on the consolidated database to refill the pool of primary key values in the key pool table.

## Results

The primary key pool table on the consolidated database is replenished and the new primary keys are replicated to the remote databases.

### 1.2.9.2.4 Using the Primary Keys from the Key Pool (SQL)

Use primary keys to form the remote user's pool of available primary keys when remote users add new customers.

## Prerequisites

The primary key pool table must already have been created.

## Context

When a sales representative adds a new customer to the Customers table, the primary key value to be inserted is obtained using a stored procedure. This example uses a stored procedure to supply the primary key value, and a stored procedure to do the insert.

## Procedure

1. Create a procedure to run on the remote databases to obtain a primary key from the primary key pool table.

For example, the NewKey procedure supplies an integer value from the key pool and deletes the value from the pool.

```
CREATE PROCEDURE NewKey(
    IN @table_name VARCHAR(40),
    OUT @value INTEGER )
BEGIN
    DECLARE NumValues INTEGER;

    SELECT COUNT(*), MIN(value)
    INTO NumValues, @value
    FROM KeyPool
    WHERE table_name = @table_name
    AND location = CURRENT PUBLISHER;
    IF NumValues > 1 THEN
        DELETE FROM KeyPool
        WHERE table_name = @table_name
        AND value = @value;
    ELSE
        // Never take the last value, because
        // ReplenishPool will not work.
        // The key pool should be kept large enough
        // that this never happens.
        SET @value = NULL;
    END IF;
END;
```

The NewKey procedure takes advantage of the fact that the Sales Representative identifier is the CURRENT PUBLISHER of the remote database.

2. Create a procedure that runs on the remote databases to insert a new row in a subscribed table.

For example, the NewCustomers procedure inserts a new customer into the table, using the value obtained by NewKey to construct the primary key.

```
CREATE PROCEDURE NewCustomers(
    IN customer_name CHAR( 40 ) )
BEGIN
    DECLARE new_cust_key INTEGER ;
    CALL NewKey( 'Customers', new_cust_key );
    INSERT
    INTO Customers (
        cust_key,
        name,
        location
    )
    VALUES (
        'Customers ' ||
        CONVERT (CHAR(3), new_cust_key),
        customer_name,
        CURRENT PUBLISHER
    );
END
```

You can enhance this procedure by testing the new\_cust\_key value obtained from NewKey to check that it is not NULL, and prevent the insert if it is NULL.

## Results

The primary key is set.

### 1.2.10 Row Partitioning Among Remote Databases

Each remote database can contain a different subset of the data stored in the consolidated database. You can create your publications and subscriptions 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**.

Sometimes, the rows of a table need to be partitioned even when the subscription expression does not exist in the table.

Sometimes, when there is a many-to-many relationship in the database, the tables need to be partitioned.

## Related Information

[Disjoint Data Partitions \[page 67\]](#)


[Overlap Partitions \[page 73\]](#)

### 1.2.11 Disjoint Data Partitions

Data partitioning is disjoint when the remote databases do not share data.

For example, each sales representative has their own set of customers and they do not share customers with other sales representatives.

In the following example, three tables store information about the interactions between sales representatives and customers: Customers, Contacts, and SalesReps. Each sales representative sells to several customers. For some customers, there is a single contact, and for other customers there are multiple contacts.



## Description of Contacts, Customers, and SalesReps Tables

The following table describes the Customers, Contacts, and SalesReps database tables.

Table	Description	Table definition
<a href="#">Contacts</a>	<p>All individual contacts that do business with the company. Each contact belongs to a single customer. The Contacts table includes the following columns:</p> <p><b>contact_key</b></p> <p>An identifier for each contact. This is the primary key.</p> <p><b>name</b></p> <p>The name of each contact.</p> <p><b>cust_key</b></p> <p>An identifier for the customer to which the contact belongs. This is a foreign key to the Customers table.</p>	<pre>CREATE TABLE Contacts (   contact_key   CHAR(12) NOT NULL,   name CHAR(40) NOT   NULL,   cust_key CHAR(12)   NOT NULL,   FOREIGN KEY   REFERENCES Customers,   PRIMARY KEY   (contact_key)   );</pre>
<a href="#">Customers</a>	<p>All customers that do business with the company. The Customers table includes the following columns:</p> <p><b>cust_key</b></p> <p>An identifier for each customer. This is the primary key.</p> <p><b>name</b></p> <p>The name of each customer.</p> <p><b>rep_key</b></p> <p>An identifier for the sales representative in a sales relationship. This is a foreign key to the SalesReps table.</p>	<pre>CREATE TABLE Customers (   cust_key CHAR(12)   NOT NULL,   name CHAR(40) NOT   NULL,   rep_key CHAR(12) NOT   NULL,   FOREIGN KEY   REFERENCES SalesReps,   PRIMARY KEY   (cust_key)   );</pre>
<a href="#">SalesReps</a>	<p>All sales representatives that work for the company. The SalesReps table includes the following columns:</p> <p><b>rep_key</b></p> <p>An identifier for each sales representative. This is the primary key.</p> <p><b>name</b></p> <p>The name of each sales representative.</p>	<pre>CREATE TABLE SalesReps (   rep_key CHAR(12) NOT   NULL,   name CHAR(40) NOT   NULL,   PRIMARY KEY (rep_key)   );</pre>

A sales representative must subscribe to a publication that provides the following information:

#### **A list of the all the sales representatives working for the company**

The following statement creates a publication that publishes the entire SalesReps table:

```
CREATE PUBLICATION SalesRepData (  
    Table SalesReps ...)  
);
```

#### **A list of customers assigned to them**

This information is available in the Customers table. The following statement creates a publication that publishes the Customers table, which contains the rows that match the value of the rep\_key column in the Customers table:

```
CREATE PUBLICATION SalesRepData (  
    TABLE Customers SUBSCRIBE BY rep_key ...  
);
```

#### **A list of the contact information for their assigned customers**

This information is available in the Contacts table. The Contacts table must be partitioned among the sales representatives, but there is no reference to the rep\_key value in the SalesReps table. To solve this problem, you can use a subquery in the Contacts article that references the rep\_key column of the Customers table.

The following statement creates a publication that publishes the Contacts table, which contains the rows that reference the rep\_key column of the Customers table.

```
CREATE PUBLICATION SalesRepData ( ...  
    TABLE Contacts  
    SUBSCRIBE BY (SELECT rep_key  
        FROM Customers  
        WHERE Contacts.cust_key = Customers.cust_key )  
);
```

One row in the Customers table has the cust\_key value in the current row of the Contacts table; the WHERE clause in the SUBSCRIBE BY statement ensures that the subquery only returns a single value.

The following statement creates the complete publication:

```
CREATE PUBLICATION SalesRepData (  
    TABLE SalesReps,  
    TABLE Customers  
    SUBSCRIBE BY rep_key,  
    TABLE Contacts  
    SUBSCRIBE BY (SELECT rep_key  
        FROM Customers  
        WHERE Contacts.cust_key = Customers.cust_key )  
);
```


#### **In this section:**

##### [BEFORE UPDATE Triggers \[page 70\]](#)

In the following example, three tables store information about the interactions between sales representatives and customers: Customers, Contacts, and SalesReps. Each sales representative sells to several customers. For some customers, there is a single contact, and for other customers there are multiple contacts.

## 1.2.11.1 BEFORE UPDATE Triggers

In the following example, three tables store information about the interactions between sales representatives and customers: Customers, Contacts, and SalesReps. Each sales representative sells to several customers. For some customers, there is a single contact, and for other customers there are multiple contacts.



A sales representative subscribes to a publication that provides a copy of the SalesReps table, a copy of the Customers table with the details of the customers assigned to them, and a copy of the Contacts table with the details of the contacts that correspond to their customers. For example, each sales representative subscribes to the following publication:

```
CREATE PUBLICATION SalesRepData (
    TABLE SalesReps,
    TABLE Customers
        SUBSCRIBE BY rep_key,
    TABLE Contacts
        SUBSCRIBE BY (SELECT rep_key
            FROM Customers
            WHERE Contacts.cust_key = Customers.cust_key )
);
```

## Maintaining Referential Integrity

This reassignment of rows among subscribers is sometimes called **territory realignment** because it is a common feature of sales force automation applications, where customers are periodically reassigned among representatives.

On the consolidated database, when a customer is reassigned to a new sales representative, the rep\_key value in the Customers table is updated.

The following statement reassigns a customer, cust1, to another sales representative, rep2.

```
UPDATE Customers
SET rep_key = 'rep2'
WHERE cust_key = 'cust1';
```

This update is replicated:

- As a DELETE statement to the Customers table on the old sales representative's remote database.
- As an INSERT statement to the Customers table on the new sales representative's remote database.

The Contacts table is *not* changed. There are no entries in the consolidated database transaction log about the Contacts table. As a result, SQL Remote on the remote databases cannot reassign the cust\_key rows of the Contacts table. This inability causes the following referential integrity problem: the Contacts table on the remote database of the old sales representative contains a cust\_key value for which there is no longer a customer.

A solution is to use a BEFORE UPDATE trigger. A BEFORE UPDATE trigger does not make any change to the database tables, but does create an entry in the consolidated database transaction log.

This BEFORE UPDATE trigger must be fired:

- Before the UPDATE statement is executed, so that the BEFORE value of the row is evaluated and added to the transaction log.
- FOR EACH ROW rather than for each statement. The information provided by the trigger must be the new subscription expression.

For example, the following statement creates a BEFORE UPDATE trigger.

```
CREATE TRIGGER "UpdateCustomer" BEFORE UPDATE OF "rep_key"
// only fire the trigger when rep_key is modified, not any other column
ORDER 1 ON "Cons"."Customers"
/* REFERENCING OLD AS old_name NEW AS new_name */
REFERENCING NEW AS NewRow
    OLD AS OldRow
FOR EACH ROW
BEGIN
// determine the new subscription expression
// for the Customers table
    UPDATE Contacts
    PUBLICATION SalesRepData
    OLD SUBSCRIBE BY ( OldRow.rep_key )
    NEW SUBSCRIBE BY ( NewRow.rep_key )
    WHERE cust_key = NewRow.cust_key;
END
;
```

SQL Remote uses the information placed in the transaction log to determine which subscribers receive which rows.

The consolidated database transaction log contains two entries after this statement is executed:

- INSERT and DELETE statements for the Contacts table generated by the BEFORE UPDATE trigger.

```
--BEGIN TRIGGER-1029-0000461705
--BEGIN TRANSACTION-1029-0000461708
BEGIN TRANSACTION
go
--UPDATE PUBLICATION-1029-0000461711 Cons.Contacts
--PUBLICATION-1029-0000461711-0002-NEW_SUBSCRIBE_BY-rep2
--PUBLICATION-1029-0000461711-0002-OLD_SUBSCRIBE_BY-rep1
--NEW-1029-0000461711
--INSERT INTO Cons.Contacts (contact_key,name,cust_key)
--VALUES ('5','Joe','cust1')
go
--OLD-1029-0000461711
--DELETE FROM Cons.Contacts
-- WHERE contact_key='5'
go
--END TRIGGER-1029-0000461743
```

- The original UPDATE statement that was executed, as well as INSERT and DELETE statements for those users that gained or lost the row respectively.

```
--PUBLICATION-1029-0000461746-0002-NEW_SUBSCRIBE_BY-rep2
--PUBLICATION-1029-0000461746-0002-OLD_SUBSCRIBE_BY-rep1
--NEW-1029-0000461746
--INSERT INTO Cons.Customers (cust_key,name,rep_key)
--VALUES ('cust1','company1','rep2')
go
--OLD-1029-0000461746
```

```

--DELETE FROM Cons.Customers
-- WHERE cust_key='cust1'
go
--UPDATE-1029-0000461746
UPDATE Cons.Customers
    SET rep_key='rep2'
VERIFY (rep_key)
VALUES ('1')
    WHERE cust_key='cust1'
go
--COMMIT-1029-0000461785
COMMIT WORK

```

SQL Remote scans the transaction log for the BEFORE and AFTER tags. Based on this information, it can determine which remote users get an INSERT, UPDATE, or DELETE statement.

- When a user is in the BEFORE list and not in the AFTER list, then a DELETE statement is sent on the Contacts table.
- When a user is in the AFTER list and not the BEFORE list, then an INSERT statement is sent on the Contacts table.
- When a user is in both the BEFORE and AFTER lists, nothing is done to the Contacts table but the UPDATE statement on the Customers table is sent.

When the BEFORE and AFTER lists are identical, the remote user already has the row and an UPDATE statement is sent.

## Notes on the Trigger

In the following example, you must use a BEFORE UPDATE trigger. In other contexts, BEFORE DELETE and BEFORE INSERT are necessary.

```

UPDATE table-name
PUBLICATION pub-name
    SUBSCRIBE BY sub-expression
WHERE search-condition;

```

In this example, you use a BEFORE trigger.

```

UPDATE table-name
PUBLICATION publication-name
    OLD SUBSCRIBE BY old-subscription-expression
    NEW SUBSCRIBE BY new-subscription-expression
WHERE search-condition;

```

The UPDATE statement lists the affected publication and table. The WHERE clause in the statement describes the affected rows. This UPDATE statement does not change the data in the table; it makes entries in the transaction log.

In this example, the subscription expression returns a single value. However, subqueries returning multiple values can also be used. The value of the subscription expression must be the value after the update.

In this example, the only subscriber to the row is the new sales representative.




## Related Information

[Overlap Partitions \[page 73\]](#)

### 1.2.12 Overlap Partitions

Data partitioning overlaps when the remote databases share data. For example, sales representatives can share customers.

Suppose three tables store information about the interactions between sales representatives and customers: Customers, Policy, and SalesReps. Each sales representative sells to several customers, and some customers deal with more than one sales representative. The Policy table has foreign keys to both the Customers and SalesReps tables. There is a many-to-many relationship between Customers and SalesReps.



### Description of Customers, Policy, and SalesReps Tables

The following table describes Customers, Policy, and SalesReps database tables.

Table	Description
<i>Customers</i>	<p>All customers that do business with the company. The Customers table has the following columns:</p> <ul style="list-style-type: none"><li><b>cust_key</b> A primary key column containing an identifier for each customer.</li><li><b>name</b> A column containing the name of each customer.</li></ul> <p>The following statements create this table:</p> <pre>CREATE TABLE Customers (   cust_key CHAR(12) NOT NULL,   name CHAR(40) NOT NULL,   PRIMARY KEY (cust_key) );</pre>

Table	Description
<a href="#">Policy</a>	<p>A three-column table that maintains the many-to-many relationship between customers and sales representatives. The Policy table has the following columns:</p> <p><b>policy_key</b></p> <p>A primary key column containing an identifier for the sales relationship.</p> <p><b>cust_key</b></p> <p>A column containing a foreign key for the customer representative in a sales relationship.</p> <p><b>rep_key</b></p> <p>A column containing a foreign key for the sales representative in a sales relationship.</p> <p>The following statements create this table:</p> <pre>CREATE TABLE Policy (   policy_key CHAR(12) NOT NULL,   cust_key CHAR(12) NOT NULL,   rep_key CHAR(12) NOT NULL,   FOREIGN KEY ( cust_key )   REFERENCES Customers ( cust_key ),   FOREIGN KEY ( rep_key )   REFERENCES SalesReps ( rep_key ),   PRIMARY KEY ( policy_key ) );</pre>
<a href="#">SalesReps</a>	<p>All sales representatives that work for the company. The SalesReps table has the following columns:</p> <p><b>rep_key</b></p> <p>An identifier for each sales representative. This is the primary key.</p> <p><b>name</b></p> <p>The name of each sales representative.</p> <p>The following statements create this table:</p> <pre>CREATE TABLE SalesReps (   rep_key CHAR(12) NOT NULL,   name CHAR(40) NOT NULL,   PRIMARY KEY (rep_key) );</pre>

## Partitioning Data

The many-to-many relationship between customers and sales representatives introduces new challenges for sharing information properly.

Sales representatives must subscribe to a publication that provides the following information:

### The entire SalesReps table

There are no qualifiers to this article, so the entire SalesReps table is included in the publication.

```
...
TABLE SalesReps,
...
```

### Those rows from the Policy table that include sales relationships involving the sales representative subscribed to the data

This article uses a SUBSCRIBE BY subscription expression to specify a column used to partition the data among the sales representatives:

```
...
TABLE Policy
SUBSCRIBE BY rep_key,
...
```

The subscription expression ensures that each sales representative receives only those rows in the table for which the value of the rep\_key column matches the value provided in the subscription.

The Policy table partitioning is **disjoint**: there are no rows that are shared with more than one subscriber.

### Those rows from the Customers table listing customers that deal with the sales representative subscribed to the data

The Customers table has no reference to the sales representative value that is used in the subscriptions to partition the data. This problem can be addressed by using a subquery in the publication.

Each row in the Customers table may be related to many rows in the SalesReps table, and shared with many sales representatives' databases. That is, there are overlapping subscriptions.

A subscription expression with a subquery is used to define the partition. The article is defined as follows:

```
...
TABLE Customers SUBSCRIBE BY (
  SELECT rep_key
  FROM Policy
  WHERE Policy.cust_key =
    Customers.cust_key
),
...
```

The Customers partitioning is **non-disjoint**: some rows are shared with more than one subscriber.

The following statement creates the complete publication:

```
CREATE PUBLICATION SalesRepData (
  TABLE SalesReps,
  TABLE Policy SUBSCRIBE BY rep_key,
  TABLE Customers SUBSCRIBE BY (
    SELECT rep_key FROM Policy
    WHERE Policy.cust_key =
      Customers.cust_key
  )
);
```

## Multiple-Valued Subqueries in Publications

The subquery in the Customers article returns a single column (rep\_key) in its result set, but may return multiple rows, corresponding to all the sales representatives that deal with the particular customer. When a subscription expression has multiple values, the row is replicated to all subscribers whose subscription matches any of the values. This ability to have multiple-valued subscription expressions allows overlapping partitioning of a table.

### In this section:

#### [Referential Integrity Maintenance When Reassigning Rows Among Subscribers \[page 76\]](#)

When rows are added, updated or deleted in one table, they may not be replicated in related tables. This scenario is handled using triggers.

#### [subscribe\\_by\\_remote Option with Many-to-many Relationships \[page 79\]](#)

When the subscribe\_by\_remote option is set to On, operations from remote databases on rows with a SUBSCRIBE BY value of NULL or an empty string assume the remote user is subscribed to the row. By default, the subscribe\_by\_remote option is set to On.

## Related Information

[CREATE PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

### 1.2.12.1 Referential Integrity Maintenance When Reassigning Rows Among Subscribers

When rows are added, updated or deleted in one table, they may not be replicated in related tables. This scenario is handled using triggers.

For example, to cancel a sales relationship between a customer and a sales representative, a row in the Policy table is deleted. In this example, the Policy table change is properly replicated to the old sales representative. However, no change has been made to the Customers table, and so no changes to the Customers table are replicated to the old sales representative.

In the absence of triggers, this can leave a subscriber with incorrect data in their Customers table. The same kind of problem arises when a new row is added to the Policy table.

## Using Triggers to Solve the Problem

The solution is to write BEFORE triggers that fire when changes are made to the Policy table. These special triggers makes no changes to the database tables, but they do make an entry in the transaction log that SQL Remote uses to maintain data in subscriber databases.

## A BEFORE INSERT Trigger

For example, the following statements create a BEFORE INSERT trigger that tracks inserts into the Policy table, and ensures that remote databases contain the proper data.

```
CREATE TRIGGER InsPolicy
BEFORE INSERT ON Policy
REFERENCING NEW AS NewRow
FOR EACH ROW
BEGIN
    UPDATE Customers
    PUBLICATION SalesRepData
    SUBSCRIBE BY (
        SELECT rep_key
        FROM Policy
        WHERE cust_key = NewRow.cust_key
        UNION ALL
        SELECT NewRow.rep_key
    )
    WHERE cust_key = NewRow.cust_key;
END;
```

## A BEFORE DELETE Trigger

The following statements create a BEFORE DELETE trigger that tracks deletes from the Policy table:

```
CREATE TRIGGER DelPolicy
BEFORE DELETE ON Policy
REFERENCING OLD AS OldRow
FOR EACH ROW
BEGIN
    UPDATE Customers
    PUBLICATION SalesRepData
    SUBSCRIBE BY (
        SELECT rep_key
        FROM Policy
        WHERE cust_key = OldRow.cust_key
        AND Policy_key <> OldRow.Policy_key
    )
    WHERE cust_key = OldRow.cust_key;
END;
```

The SUBSCRIBE BY clause of the UPDATE PUBLICATION statement contains a subquery, and this subquery can be multiple-valued.

## Multiple-Valued Subqueries

The subquery in the SUBSCRIBE clause of the UPDATE PUBLICATION is a UNION expression, and can be multiple-valued:

```
...
SELECT rep_key
FROM Policy
WHERE cust_key = NewRow.cust_key
```

```
UNION ALL
SELECT NewRow.rep_key
...
```

- The first part of the UNION is the set of existing sales representatives dealing with the customer, taken from the Policy table.  
The result set of the subscription query must be all those sales representatives receiving the row, not just the new sales representatives.
- The second part of the UNION is the rep\_key value for the new sales representative dealing with the customer, taken from the INSERT statement.

The subquery in the BEFORE DELETE trigger is multiple-valued:

```
...
SELECT rep_key
FROM Policy
WHERE cust_key = OldRow.cust_key
AND rep_key <> OldRow.rep_key
...
```

- The subquery takes rep\_key values from the Policy table. The values include the primary key values of all those sales representatives who deal with the customer being transferred (WHERE *cust\_key* = *OldRow.cust\_key*), with the exception of the one being deleted (AND *rep\_key* <> *OldRow.rep\_key*).  
The result set of the subscription query must be all those values matched by sales representatives receiving the row following the delete.

## Notes

- Data in the Customers table is not identified with an individual subscriber (by a primary key value, for example) and is shared among more than one subscriber. This allows the possibility of the data being updated at more than one remote site between replication messages, which can lead to replication conflicts. You can address this issue either by privileges (allowing only certain users the right to update the Customers table, for example) or by adding RESOLVE UPDATE triggers to the database to handle the conflicts programmatically.
- Updates on the Policy table have not been described here. Either they should be prevented, or a BEFORE UPDATE trigger must be created that combines features of the BEFORE INSERT and BEFORE DELETE triggers shown in the example.


## 1.2.12.2 subscribe\_by\_remote Option with Many-to-many Relationships

When the `subscribe_by_remote` option is set to On, operations from remote databases on rows with a `SUBSCRIBE BY` value of NULL or an empty string assume the remote user is subscribed to the row. By default, the `subscribe_by_remote` option is set to On.

The `subscribe_by_remote` option solves a problem that otherwise would arise with some publications. The following publication uses a subquery for the Customers table subscription expression because customers can belong to several sales representatives:

```
CREATE PUBLICATION SalesRepData (  
  TABLE SalesReps,  
  TABLE Policy SUBSCRIBE BY rep_key,  
  TABLE Customers SUBSCRIBE BY (  
    SELECT rep_key FROM Policy  
    WHERE Policy.cust_key =  
      Customers.cust_key  
  ),  
);
```

For example, Marc Dill is a Sales representative who has just arranged a policy with a new customer. He inserts a new row in the Customers table and inserts a row in the Policy table to assign the new customer to himself.



On the consolidated database, SQL Remote carries out the insert of the Customers row and SQL Anywhere records the subscription value in the transaction log, at the time of the insert.

Later, when the SQL Remote scans the transaction log, it builds a list of subscribers from the subscription expression, and Marc Dill is not on the list, as the row in the Policy table assigning the customer to him has not yet been applied. If `subscribe_by_remote` were set to Off, the result would be that the new Customer is sent back to Marc Dill as a DELETE statement.

As long as `subscribe_by_remote` is set to On, the SQL Remote assumes the row belongs to the Sales representative who inserted it, the INSERT statement is not replicated back to Marc Dill, and the replication system is intact.

If `subscribe_by_remote` is set to Off, you must ensure that the Policy row is inserted before the Customers row, with the referential integrity violation avoided by postponing checking to the end of the transaction.

### Related Information

[subscribe\\_by\\_remote Option \[SQL Remote\]](#)

## 1.2.13 Unique Identification Numbers for Remote Databases

SQL Remote requires that you assign a different identification number to each remote 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 user name.

### Using the `global_database_id` Option

The public option `global_database_id` in each database must be set to a unique, non-negative integer. The range of default values for a particular database is  $p_n + 1$  to  $p(n + 1)$ , where  $p$  is the partition size and  $n$  is the value of the public option `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.

When `global_database_id` is set to a non-negative integer, SQL Anywhere chooses default values by applying the following rules:

- When the column contains no values in the current partition, the first default value is  $p_n + 1$ .
- When 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  $p_n + 1$  or greater than  $p(n + 1)$ . Such values may be present if they have been replicated from another database via MobiLink synchronization.

If the public option `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, when the column is contained in the table's primary key.

Because the public option `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 example, a new value of `global_database_id` 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 when the column does not permit nulls. To detect that the supply of unused values is low and handle this condition, create an event of type `GlobalAutoincrement`.

Should the values in a particular partition become exhausted, you can assign a new 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 do some other action) when the partition is nearly exhausted.

#### In this section:

[Setting the `global\_database\_id` Value \(SQL\) \[page 81\]](#)

Set an identification number for remote databases.



[Setting Database Identification Numbers When Extracting Databases \(SQL\) \[page 82\]](#)

Write a stored procedure to automate the task of setting unique database identification numbers.

## Related Information

[Primary Key Pools with SQL Remote \[page 61\]](#)

[DEFAULT GLOBAL AUTOINCREMENT Declaration \[page 60\]](#)

[Trigger Conditions for Events](#)

[global\\_database\\_id Option](#)

### 1.2.13.1 Setting the global\_database\_id Value (SQL)

Set an identification number for remote databases.

#### Context

The identification number must be a non-negative integer.

#### Procedure

Set the value of the global\_database\_id option.

For example, 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.

#### Results

The global\_database\_id number is set.

## Related Information

[global\\_database\\_id Option](#)

## 1.2.13.2 Setting Database Identification Numbers When Extracting Databases (SQL)

Write a stored procedure to automate the task of setting unique database identification numbers.

### Prerequisites

Any user with the MANAGE REPLICATION system privilege can create hook procedures. However, to ensure that the hook can access the #hook\_dict table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the SELECT ANY TABLE and UPDATE ANY TABLE system privileges.
- Be defined using the SQL SECURITY INVOKER clause of the CREATE PROCEDURE statement.

### Procedure

1. Create a stored procedure named `sp_hook_dbxtract_begin`.

For example, to extract a database for remote user user2 with a user\_id of 1001, execute the following statements:

```
SET OPTION "PUBLIC"."global_database_id" = '1';
CREATE TABLE extract_id (next_id INTEGER NOT NULL) ;
INSERT INTO extract_id VALUES( 1 );
CREATE PROCEDURE sp_hook_dbxtract_begin
AS
    DECLARE @next_id  INTEGER
    UPDATE extract_id SET next_id = next_id + 1000
    SELECT @next_id = (next_id)
    FROM extract_id
    COMMIT
    UPDATE #hook_dict
    SET VALUE = @next_id
    WHERE NAME = 'extracted_db_global_id';
```

Each extracted or re-extracted database gets a different global\_database\_id. The first starts at 1001, the next at 2001, and so on.

2. Run the Extraction utility (dbxtract) with the -v option or the [Extract Database Wizard](#) to extract your remote databases. The Extraction utility performs the following tasks:
  - a. Creates a temporary table name #hook\_dict, with the following contents:

name	value
extracted_db_global_id	The user ID being extracted.

When you write an sp\_hook\_dbxtract\_begin procedure to modify the value column of the row, that value is used as the global\_database\_id option of the extracted database, and marks the beginning of the range of primary key values for DEFAULT GLOBAL AUTOINCREMENT values.

- When you do not define an `sp_hook_dbxtract_begin` procedure, the extracted database has a `global_database_id` set to 101.
  - When you define an `sp_hook_dbxtract_begin` procedure that does not modify any rows in the `#hook_dict`, then the `global_database_id` is still set to 101.
- b. Calls `sp_hook_dbxtract_begin`.
  - c. Outputs the following information to assist in debugging procedure hooks:
    - The procedure hooks found.
    - The contents of `#hook_dict` before the procedure hook is called.
    - The contents of `#hook_dict` after the procedure hook is called.

## Results

The unique database identification numbers are set.

## Related Information

[The #hook\\_dict Table \[page 242\]](#)

[SQL Remote System Procedures \[page 241\]](#)

[global\\_database\\_id Option](#)

[Extraction Utility \(dbxtract\) \[page 225\]](#)

## 1.3 SQL Remote Security Considerations

SQL Remote provides a number of features to assist in protecting your data.

### **SYS\_RUN\_REPLICATION\_ROLE System Role**

Connect to the SQL Remote Message Agent (dbremote) with a user that has the `SYS_RUN_REPLICATION_ROLE` system role. A user granted the `SYS_RUN_REPLICATION_ROLE` system role cannot exercise the privileges inherited from this role on any connection apart from SQL Remote or `dbmsync`. Therefore, even if the user name and password for a user with the `SYS_RUN_REPLICATION_ROLE` system role are widely distributed, security is not compromised. As long as the user name has no privileges beyond `CONNECT` granted on the database, no one can use this user name to access data in the database.

### **Database Encryption**

You can encrypt your database using the `-ea` and `-ek` options of the extraction utility, `dbxtract`. You can also encrypt data at the table level, using the `dbxtract -et` option.

#### **-ea**

Specifies the encryption algorithm for the new database. This option allows you to choose a strong encryption algorithm to encrypt your new database. You can choose either AES (the default) or

AES\_FIPS for the FIPS-certified algorithm. AES\_FIPS uses a separate library and is not compatible with AES.

For greater security, specify AES or AES256 for 128-bit or 256-bit strong encryption, respectively. Specify AES\_FIPS or AES256\_FIPS for 128-bit or 256-bit FIPS-certified encryption, respectively. For strong encryption, you must also specify the -ek or -ep option.

#### **-ek**

Specifies the encryption key for the new database. This option allows you to create a strongly encrypted database by specifying an encryption key directly in the command. The algorithm used to encrypt the database is AES or AES\_FIPS as specified by the -ea option. If you specify the -ek option without specifying -ea, the AES algorithm is used.

#### **Caution**

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

#### **-et**

Enables database table encryption in a new database. If you specify the -et option without the -ea option, the AES algorithm is used. If you specify the -et option, you must also specify -ep or -ek. You can change the table encryption settings for a new database to be different from those of the database you are unloading.

### **Message Encryption**

The SQL Remote Message Agent (dbremote) uses a simple obfuscation algorithm to protect messages against casual snooping. However, this encryption scheme is not intended to provide full protection against determined efforts to decipher them. You can implement your own custom encoding scheme if security of your data over the network is an issue.

## **1.4 Management of SQL Remote Systems**

You deploy and administer a SQL Remote system from the consolidated database.

To administer SQL Remote systems You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

To deploy and administer a SQL Remote system:

1. Set up the consolidated database.
2. Review and test your SQL Remote system.  
Thorough testing of your SQL Remote system should be done before deployment, especially if you have a large number of remote databases.
3. Create remote databases and deploying the design.  
As the DBA of the consolidated database, you deploy SQL Remote by:
  1. Creating a SQL Anywhere database for each remote user, with their own initial copy of the data, and starting their subscriptions.
  2. Installing on each remote user's computer the SQL Anywhere database server, the remote database, SQL Remote, and the client application.

4. Run the SQL Remote Message Agent (dbremote) to exchange messages.  
To exchange messages, you need to:
  1. Decide whether to run the SQL Remote Message Agent (dbremote) in continuous mode or batch mode on the consolidated and remote databases.
  2. Ensure that the system is properly configured with correct user names, SQL Remote Message Agent (dbremote) connection strings, privileges, and so on.
5. Manage messages.  
Use the Guaranteed Message Delivery System to manage the messages being sent back and forth among many databases.
6. Improve performance.
7. Implement a backup and recovery strategy.  
You must create and implement a backup and recovery strategy for the consolidated database.
8. Handle errors.
9. Upgrade the software and database schemas as required.

#### **In this section:**

##### [Remote Database Extraction \[page 86\]](#)

To create a database for a remote user, you **extract** the remote database from the consolidated database.

##### [Remote Database Extraction to a Script File \[page 88\]](#)

In most deployment scenarios, you must customize the extraction and creation of remote databases. You can create a custom extraction by choosing to extract the database into a script file and then modifying the script.

##### [SQL Remote Message Agent \(dbremote\) \[page 95\]](#)

The SQL Remote Message Agent (dbremote) is a key component in SQL Remote replication. It must be installed and run on every database in the system.

##### [SQL Remote Performance \[page 102\]](#)

SQL Remote divides the workload for notifications on subscriptions and publications between the database server and the SQL Remote Message Agent (dbremote) to manage network performance.

##### [Guaranteed Message Delivery System \[page 114\]](#)

The Guaranteed Message Delivery System ensures that all replicated operations are applied in the correct order, no replicated operations are missed, and no replicated operation is applied twice.

##### [Message Size \[page 118\]](#)

The SQL Remote Message Agent (dbremote) provides a number of encoding and compression features.

##### [SQL Remote Message Systems \[page 120\]](#)

In SQL Remote replication, a message system is a protocol for exchanging messages between the consolidated database and a remote database. SQL Remote exchanges data among databases using one or more underlying message systems.

##### [SQL Remote System Backups \[page 138\]](#)

There are a number of considerations to be aware of before you perform backup activities.

##### [Recovering a Consolidated Database Manually \(Command Line\) \[page 145\]](#)

Recover a consolidated database by applying each transaction log to the database.

##### [Recovering a Consolidated Database Automatically \(Command Line\) \[page 147\]](#)

Recover a consolidated database automatically.

[Replication Error Reporting and Handling \[page 149\]](#)

SQL Remote provides a number of methods for handling potential errors.

[SQL Remote Security Considerations \[page 154\]](#)

SQL Remote provides a number of features to protect your data.

[Upgrades and Resynchronization \[page 155\]](#)

There are a number of different ways to upgrade and resynchronize SQL Remote databases.

[SQL Remote Passthrough Mode \[page 157\]](#)

Use passthrough mode to pass standard SQL statements to a remote database where they can be executed.

[Subscription Resynchronization \[page 160\]](#)

When you create a remote database, you extract both the schema and data from the consolidated database and use them to build the remote database. This process ensures that each database has an initial copy of the data.

## Related Information

[Create a SQL Remote System \[page 11\]](#)

[SQL Remote Message Agent \(dbremote\) Modes \[page 96\]](#)

[Embedded Database Application Deployment](#)

[SQL Remote Deployment](#)

[GRANT ROLE Statement](#)

### 1.4.1 Remote Database Extraction

To create a database for a remote user, you **extract** the remote database from the consolidated database.

You can use either the *Extract Database Wizard* or the Extraction utility (dbxtract) to extract a remote database from a consolidated database for a specified remote user. Either method allows you to do one or more of the following tasks:

#### **Automatically extract and reload the schema and data directly into a new or existing database**

This is an appropriate method to use when learning about SQL Remote. If you use this method, no interim copy of the data is created on disk. This method provides greater security for your data. However, it is more time consuming to implement.

#### **Extract the schema and data to files, and then load them into a new or existing database**

When deploying SQL Remote, this method is preferred. You can edit the schema file to customize the extraction and creation of your remote databases.

One method to increase efficiency is to create more than one remote database.

#### **In this section:**

[Extracting a Remote Database Automatically \(SQL Central\) \[page 87\]](#)

Extract a consolidated database and reload the schema and data into a new database. No interim copy of the data is created on disk.

## Related Information

[Remote Database Extraction to a Script File \[page 88\]](#)

[Creating Multiple Remote Databases \(Command Line\) \[page 93\]](#)

### 1.4.1.1 Extracting a Remote Database Automatically (SQL Central)

Extract a consolidated database and reload the schema and data into a new database. No interim copy of the data is created on disk.

## Prerequisites

You must have the EXECUTE ANY PROCEDURE and SELECT ANY TABLE privileges. The SELECT ANY TABLE privilege comes with the SYS\_REPLICATION\_ADMIN\_ROLE role.

## Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the consolidated database.
2. Click **Tools** > *SQL Anywhere 17* > *Extract Database*.
3. Choose the consolidated database to connect to and extract.
4. When prompted, click *Extract And Reload Into A New Database*.

When prompted, click *Extract Structure And Data*.

5. Follow the instructions in the wizard and accept the defaults.

## Results

The new remote database is created with the appropriate schema, remote users, publications, subscriptions, and triggers. By default, the data from the consolidated database is extracted to the remote database and the subscriptions are started. However, the wizard does not start the SQL Remote Message Agent, so no messages are exchanged.

## Related Information

[SQL Remote Message Agent \(dbremote\) \[page 95\]](#)

[Remote Database Extraction to a Script File \[page 88\]](#)

[Extraction Utility \(dbxtract\) \[page 225\]](#)

### 1.4.2 Remote Database Extraction to a Script File

In most deployment scenarios, you must customize the extraction and creation of remote databases. You can create a custom extraction by choosing to extract the database into a script file and then modifying the script.

Use the Extraction utility (dbxtract) to extract the schema, remote users, publications, subscriptions, and triggers from the consolidated database. The following is a typical example:

```
dbxtract -c "UID=DBA;PWD=sql" -xv -up -r remote.sql remote2 remote_user
```

The `-xv` option causes views to be excluded, including materialized views. Materialized views that reference tables that are not being published will result in errors. There are other options to exclude certain types of database objects.

The `-up` option causes hashed passwords to be included in the script file for any users that are to be created in the remote database.

The `-r` option is used to specify the location and name for the SQL script file. In this example, `remote.sql` is the name of the script file. If this option is not used, the SQL script is written to `reload.sql` in the current directory.

The `remote2` directory is a subdirectory relative to the current working directory of the database server. By default, the data for each table extracted is unloaded into separate files and placed in this directory. Use the `-n` option if you are not interested in populating the extracted tables with data. In this case, the `remote2` directory is ignored but it must be specified.

Finally, `remote_user` is the name of a user that has been granted the REMOTE privilege.

### The Extraction Utility Script File

The Extraction utility script file contains the SQL statements necessary to build the database schema, including statements to create:

- Publishers, remote, and consolidated users
- Publications and subscriptions
- Message types
- Tables
- Views
- Triggers
- Procedures



The `-r` option is used to specify the location and name for the script file. The default name of this file is `reload.sql`.

### **i** Note

You may need to edit the script file when you create remote databases. The Extraction utility (`dbxtract`) is intended to assist in preparing remote databases, but is not intended as a black box solution for all circumstances.

#### **In this section:**

##### [Creating a Remote Database from an Existing Database \(Command Prompt\) \[page 89\]](#)

Create a remote database using an existing database as a model. A schema, remote users, publications, subscriptions, and triggers are also generated.

##### [When to Edit the `reload.sql` File \[page 91\]](#)

Edit the `reload.sql` script file as needed when creating remote databases.

##### [Database Extraction for a Multi-tier Hierarchy System \[page 92\]](#)

This example describes how to use the Extraction utility (`dbxtract`) and the *Extract Database Wizard* to create databases within a multi-tiered hierarchy.

##### [Creating Multiple Remote Databases \(Command Line\) \[page 93\]](#)

Streamline the creation of multiple databases at a time.

## **Related Information**

[SQL Remote Message Agent \(`dbremote`\) \[page 95\]](#)

[Remote Database Extraction \[page 86\]](#)

[Extracting a Remote Database Automatically \(SQL Central\) \[page 87\]](#)

[Extraction Utility \(`dbxtract`\) \[page 225\]](#)

### **1.4.2.1 Creating a Remote Database from an Existing Database (Command Prompt)**

Create a remote database using an existing database as a model. A schema, remote users, publications, subscriptions, and triggers are also generated.

#### **Prerequisites**

The steps below create a remote database using the sample database `demo.db` as a consolidated database example.

It is assumed that you know how to start and connect to the sample database.

## Procedure

1. Start and connect to the sample database using Interactive SQL, then execute the following statements.

```
CREATE REMOTE MESSAGE TYPE FILE ADDRESS 'messages\host';
GRANT PUBLISH TO DBA;
GRANT CONNECT TO remote_user IDENTIFIED BY 'sqlremotepw';
GRANT REMOTE TO remote_user TYPE FILE ADDRESS 'messages\remote' SEND EVERY
'00:01:00';
CREATE PUBLICATION PubDemo (
    TABLE Customers,
    TABLE Employees
);
CREATE SUBSCRIPTION TO PubDemo FOR remote_user;
```

These statements prepare the sample database as a consolidated database for SQL Remote replication.

2. At a command prompt, use the Extraction utility (dbxtract) to extract the database schema and data to a remote database as follows:

```
dbxtract -c "UID=DBA;PWD=sql" -an remote2\remote.db -ae -xv -up remote2
remote_user
```

The `-an` option is used to specify the location and name for the remote database. If this option is not used, then no database is created and the SQL statements required to set up the remote database are recorded in the `reload.sql` file.

The `-ae` option is used to resume processing after any error, rather than quit. Any SQL statements that fail are recorded in the `unprocessed.sql` file. In this example, there should be no failures. The `-ae` option is specified when the `-an` or `-ac` option is used.

The `-xv` option causes views to be excluded, including materialized views. Materialized views that reference tables that are not being published will result in errors. The sample database has such materialized views.

The `-up` option causes passwords to be included for any users that are created in the remote database `remote.db`. This option is not required when the `-an` or `-ac` option is used.

The `remote2` directory is a subdirectory relative to the current working directory of the database server. The directory name need not be specified when the `-an` or `-ac` option is used.

3. Start and connect to the `remote.db` database using Interactive SQL.

Use this session to verify that the Customers and Employees tables are present in the remote database.

## Results

The remote database is created.

Repeat step 2 omitting the `-an remote2\remote.db` to generate a `reload.sql` file. This file can be used to generate the schema, remote users, publications, subscriptions, and triggers of the remote database. It can be edited to fine-tune the contents of the remote database.

## Related Information

[Extraction Utility \(dbxtract\) \[page 225\]](#)

[Initialization Utility \(dbinit\)](#)

### 1.4.2.2 When to Edit the `reload.sql` File

Edit the `reload.sql` script file as needed when creating remote databases.

For example, you must edit the `reload.sql` file in the following cases:

#### Adding Unreplicated Tables to Remote Database

Remote databases can have tables that are not present at their consolidated database as long as these tables do not take part in replication. The Extraction utility (dbxtract) and *Extract Database Wizard* cannot extract unreplicated tables from a consolidated database.

After extracting the database, edit `reload.sql` to add such tables.

#### Extracting Procedures, Triggers, and Views

By default, the Extraction utility (dbxtract) and *Extract Database Wizard* extract all stored procedures, triggers, and views from the database. While some of these views and procedures are likely to be required on the remote site, others may not be required. For example, a procedure could refer to parts of the database that are not included in the remote site.

After extracting the database, edit `reload.sql` to remove unnecessary procedures, triggers, and views.

#### Using the Extraction Utility (dbxtract) in Multi-Tiered Systems

You use the Extraction utility (dbxtract) or the *Extract Database Wizard* to create second- to n-level databases within a multi-tiered system.

By default, dbxtract and *Extract Database Wizard* extract all stored procedures, triggers, and views from the database. While some of these views and procedures are likely to be required on the remote site, others may not be required. For example, a procedure could refer to parts of the database that are not included in the remote site.

After extracting the database, edit `reload.sql` to remove unnecessary procedures, triggers, and views.

## Related Information


[Creating Multiple Remote Databases \(Command Line\) \[page 93\]](#)

[Database Extraction for a Multi-tier Hierarchy System \[page 92\]](#)

### 1.4.2.3 Database Extraction for a Multi-tier Hierarchy System

This example describes how to use the Extraction utility (dbxtract) and the *Extract Database Wizard* to create databases within a multi-tiered hierarchy.

To understand the role of the Extraction utility (dbxtract) and the *Extract Database Wizard* in multi-tiered arrangements, consider a three-tiered SQL Remote system. This system is illustrated in the following diagram.



To create the remote databases for a three-tiered system:

1. Use the Extraction utility (dbxtract) on the top-level, consolidated database, HQ, to create the second-level databases Region 1 and Region 2.
2. Use the Extraction utility (dbxtract), on the second-level databases, Region 1 and Region 2, to create the third-level databases for users Laptop 1, Laptop 2, and Laptop 3. The second-level databases are remote databases to the first-level database, HQ, and are consolidated databases to the third-level databases, Laptop 1, Laptop 2, and Laptop 3.

### Re-extracting Databases in a Multi-Tier Hierarchy System

If you have to re-extract the *schema* for the second-level database from the top-level consolidated database, the Extraction utility (dbxtract) deletes the remote users (Laptop 1, Laptop 2, and Laptop 3) along with their subscriptions and privileges. As a result, you must recreate those third-level users and their subscriptions manually.

If you only have to re-extract the *data* from the second-level databases from the top-level consolidated database, the Extraction utility (dbxtract) does not affect the remote users.

## Fully Qualified Publication Definitions

Fully qualified publication definitions contain WHERE and SUBSCRIBE BY clauses. Usually you do not need to extract fully qualified publication definitions for a remote database, since the remote database typically replicates all rows back to the consolidated database.

## Related Information

[Creating Multiple Remote Databases \(Command Line\) \[page 93\]](#)

[Extraction Utility \(dbxtract\) \[page 225\]](#)

### 1.4.2.4 Creating Multiple Remote Databases (Command Line)

Streamline the creation of multiple databases at a time.

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Procedure

1. Make a copy of the consolidated database and start the subscriptions for the remote users from the consolidated database. For example:
  - a. Shut down the consolidated database and the SQL Remote Message Agent (if it is running).
  - b. Start a local copy of the consolidated database using dbeng17 and a different server name to ensure that no other processes connect to the local copy.
  - c. Start the subscriptions.

The subscriptions must be started at the same time that the consolidated database copy is made. Any operations that take place between copying the database and starting the subscriptions can be lost, and can lead to errors at remote databases. Starting subscriptions on the consolidated database allows messages to be packaged and sent to subscribers, even if the subscriber databases do not exist yet.

To start several subscriptions within a single transaction, use the REMOTE RESET statement.

- d. Immediately shut down the consolidated database.
- e. Copy the consolidated database.

By default, both the Extraction utility (dbxtract) and the *Extract Database Wizard* run at isolation level 3. This isolation level ensures that data in the extracted database is consistent with data on the

database server; however, it can prevent other users from using the database. Extract your remote database against a copy of the consolidated database.

- f. Re-start the consolidated database, and if it was running, re-start the SQL Remote Message Agent on the consolidated database.
2. Extract the remote database schema from the copy of the consolidated database. As the database is a copy, there are no locking and concurrency problems; nevertheless, for a large number of remote databases, this process can take a while.

When extracting the remote database schema, choose the following options:

- a. Extract only the schema for the remote database.

By default, both the Extraction utility (dbxtract) and the *Extract Database Wizard* extract one database at a time, including the schema and data for each user. However, in most deployment scenarios the remote databases use the same schema but different data. Using the Extraction utility (dbxtract) or the *Extract Database Wizard* to extract both schema and data for each user results in repeatedly extracting the same schema.

- b. Order the data by primary key.

By default, the data in each table is ordered by primary key. Loading data into the remote database is faster when the data is ordered by primary key using the -u option of the Extraction utility (dbxtract).

3. Create an empty remote database using the `reload.sql` file. Copy this database file to create the required number of remote databases.
4. For each remote database, define the SQL Remote definitions specific to each remote user.
5. For each remote user, extract only their corresponding data from the consolidated database using the -d option of the Extraction utility (dbxtract).
6. Load the data for each remote user into the corresponding remote database.

As each remote database is created, its data is out of date with the live consolidated database.

However, when you run the SQL Remote Message Agent (dbremote), each user can receive and apply messages that have been sent from the live consolidated database to bring themselves up to date.

## Results

The remote databases are created.

## Related Information

[When to Edit the reload.sql File \[page 91\]](#)

[SQL Remote Message Agent \(dbremote\) \[page 95\]](#)

[User Privileges \[page 22\]](#)

[Starting Subscriptions \(SQL Central\) \[page 164\]](#)

[Database Extraction for a Multi-tier Hierarchy System \[page 92\]](#)

[START SUBSCRIPTION Statement \[SQL Remote\]](#)

[REMOTE RESET Statement \[SQL Remote\]](#)

## 1.4.3 SQL Remote Message Agent (dbremote)

The SQL Remote Message Agent (dbremote) is a key component in SQL Remote replication. It must be installed and run on every database in the system.

The SQL Remote Message Agent (dbremote) handles both sending and receiving messages.

To run dbremote you must have the SYS\_RUN\_REPLICATION\_ROLE system role.

It carries out the following functions:

### SQL Remote Message Agent (dbremote) tasks when sending messages

- It scans the transaction log at each publisher database and translates the transaction log entries into messages for subscribers.
- It sends the messages to subscribers.
- When it receives a request to resend messages, the SQL Remote Message Agent (dbremote) resends the messages to the database that made the request.
- It maintains message information in the system tables, and manages the Guaranteed Message Delivery System.

### SQL Remote Message Agent (dbremote) tasks when receiving messages

- It processes incoming messages, and applies them in the proper order to the database.
- It requests that missing messages be re-sent.
- It maintains the message information in the system tables, and manages the Guaranteed Message Delivery System.

## Connections

The SQL Remote Message Agent (dbremote) uses several connections to the database server. These are:

### One global connection

This connection is active all the time the SQL Remote Message Agent (dbremote) is running.

### One connection for scanning the transaction log

This connection is active during the scan phase only.

### One connection for executing commands from the transaction log-scanning thread

This connection is active during the scan phase only.

### One connection for processing synchronize subscription requests

This connection is active during the send phase only.

### One connection for each worker thread

These connections are alive during the receive phase only.

**In this section:**

[SQL Remote Message Agent \(dbremote\) Modes \[page 96\]](#)

The SQL Remote Message Agent (dbremote) can be run in either continuous mode or batch mode.

[Running the SQL Remote Message Agent in Continuous Mode \(dbremote\) \[page 97\]](#)

Send messages at times specified by the SEND AT or SEND EVERY frequency (in continuous mode) by setting the properties of each remote user.

[Running the SQL Remote Message Agent in Batch Mode \(dbremote\) \[page 101\]](#)

Receive and process incoming messages, scan the transaction log, and create and send outgoing messages using the SQL Remote Message Agent (dbremote).

[Run the SQL Remote Message Agent \(dbremote\) on UNIX/Linux \[page 102\]](#)

On UNIX and Linux platforms, you run the SQL Remote Message Agent (dbremote) as a daemon by supplying the -ud option.

## Related Information

[Tasks to Send Messages \[page 110\]](#)

[Tasks to Receive Messages \[page 103\]](#)

### 1.4.3.1 SQL Remote Message Agent (dbremote) Modes

The SQL Remote Message Agent (dbremote) can be run in either continuous mode or batch mode.

#### Continuous mode

In continuous mode, the SQL Remote Message Agent (dbremote) periodically sends messages, at times specified by the send frequency properties of each remote user. When it is not sending messages, it receives messages as they arrive.

Continuous mode is useful at consolidated databases, where messages may be coming in and going out at any time, to spread out the workload and to ensure prompt replication.

#### Batch mode

In batch mode, the SQL Remote Message Agent (dbremote) receives and processes incoming messages, scans the transaction log once, creates and sends the outgoing messages, and then stops.

Batch mode is useful at occasionally connected remote databases, where messages can only be exchanged with the consolidated database when a connection is made, for example, when the remote database dials up to the main network.

**In this section:**

[SQL Remote Message Agent \(dbremote\) Requirements \[page 97\]](#)

SQL Remote is very flexible. Within a system, you can run the SQL Remote Message Agent (dbremote) in both modes, on multiple devices, and on multiple operating systems.



## Related Information

[Running the SQL Remote Message Agent in Continuous Mode \(dbremote\) \[page 97\]](#)

[Running the SQL Remote Message Agent in Batch Mode \(dbremote\) \[page 101\]](#)

[GRANT ROLE Statement](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

### 1.4.3.1.1 SQL Remote Message Agent (dbremote) Requirements

SQL Remote is very flexible. Within a system, you can run the SQL Remote Message Agent (dbremote) in both modes, on multiple devices, and on multiple operating systems.

However, SQL Remote has the following requirements:

#### **SYS\_RUN\_REPLICATION\_ROLE system role**

The SQL Remote Message Agent (dbremote) must be run by a user with the SYS\_RUN\_REPLICATION\_ROLE system role.

#### **The maximum message length must be the same for each SQL Remote Message Agent (dbremote) in the system**

This length can be restricted by operating system memory allocation limits. Received messages that are longer than the limit are deleted as corrupt messages. The default value is 50000 bytes. This length is configurable, using the SQL Remote Message Agent (dbremote) -l option.

### 1.4.3.2 Running the SQL Remote Message Agent in Continuous Mode (dbremote)

Send messages at times specified by the SEND AT or SEND EVERY frequency (in continuous mode) by setting the properties of each remote user.

## Prerequisites

You must have the SYS\_RUN\_REPLICATION\_ROLE system role.

To run the SQL Remote Message Agent (dbremote) in continuous mode, for example on the consolidated database, you must ensure that every REMOTE user has a send frequency specified.

The maximum message length, as defined by the -l option, must be the same for all databases in the system.

## Context

Typically, the consolidated database is run in continuous mode. In continuous mode, the SQL Remote Message Agent (dbremote) sends messages at the times specified with the SEND AT or SEND EVERY property.

## Procedure

1. Ensure that every REMOTE user either has a SEND AT or SEND EVERY frequency specified.
2. Start the SQL Remote Message Agent (dbremote) *without* the -b option.

On Windows, the SQL Remote Message Agent (dbremote) is named `dbremote.exe`. On Unix, the name is `dbremote`. On macOS, you can also use `SyncConsole` to start the SQL Remote Message Agent (dbremote).

For example, the following command runs dbremote in continuous mode on a database file named `c:\mydata.db`, connecting with user name `ManagerSteve` and password `passwd`:

```
dbremote -c "UID=ManagerSteve;PWD=passwd;DBF=c:\mydata.db" -l 40000
```

## Results

The SQL Remote Message Agent (dbremote) is now set to run in continuous mode.

### In this section:

#### [Send Frequency \[page 99\]](#)

To run the SQL Remote Message Agent (dbremote) in continuous mode, for example on the consolidated database, ensure that every REMOTE user has a specified send frequency.

#### [Setting the Send Frequency for Messages \[page 100\]](#)

Specify a send frequency for every REMOTE user to run the SQL Remote Message Agent (dbremote) in continuous mode, for example on the consolidated database.

## Related Information

[Service Utility \(dbsvc\) for Windows](#)

## 1.4.3.2.1 Send Frequency

To run the SQL Remote Message Agent (dbremote) in continuous mode, for example on the consolidated database, ensure that every REMOTE user has a specified send frequency.

In continuous mode, the SQL Remote Message Agent (dbremote) sends messages at the times specified with the SEND AT or SEND EVERY property.

The SQL Remote Message Agent (dbremote) supports the following send frequency values:

### **SEND EVERY**

Specifies a length of time to wait between sending messages.

When any user with SEND EVERY set is sent messages, all users with the same frequency are sent messages at the same time. For example, all remote users who receive updates every twelve hours are sent updates at the same time, rather than being staggered. This reduces the number of times the SQL Anywhere transaction log has to be processed. Use as few unique frequencies as possible.

A send frequency can be specified in hours, minutes, and seconds in the format HH:MM:SS.

### **SEND AT**

Specifies a time of day at which messages are sent.

Updates are sent daily at the specified time. Use as few distinct times as possible rather than staggering the send times and choose times when the database is not busy.

### **Default setting (no SEND clause)**

If any user has no SEND AT or SEND EVERY clause specified, the SQL Remote Message Agent (dbremote) runs in batch mode, sending messages every time it is run, and then stopping.

## Sending messages too frequently

If you send messages frequently, there is a greater chance of small messages being sent. Sending messages less frequently allows more instructions to be grouped in a single message. If a large number of small messages is a concern for your message system, then avoid using very small send frequency periods.

## Related Information

[Running the SQL Remote Message Agent in Batch Mode \(dbremote\) \[page 101\]](#)

[GRANT REMOTE Statement \[SQL Remote\]](#)

## 1.4.3.2.2 Setting the Send Frequency for Messages

Specify a send frequency for every REMOTE user to run the SQL Remote Message Agent (dbremote) in continuous mode, for example on the consolidated database.

### Prerequisites

You must have the SYS\_RUN\_REPLICATION\_ROLE system role.

### Context

In continuous mode, the SQL Remote Message Agent (dbremote) sends messages at the times specified with the SEND AT or SEND EVERY property.

### Procedure

1. In SQL Central, use the [SQL Anywhere 17](#) plug-in to connect to the database.
2. Double-click [SQL Remote Users](#).
3. Right-click a user and click [Properties](#).
4. Click the [SQL Remote](#) tab.
5. Click either [Send Every](#) or [Send Daily At](#) and specify a time. Click [OK](#).

### Results

The frequency is set.

### Related Information

[GRANT REMOTE Statement \[SQL Remote\]](#)

### 1.4.3.3 Running the SQL Remote Message Agent in Batch Mode (dbremote)

Receive and process incoming messages, scan the transaction log, and create and send outgoing messages using the SQL Remote Message Agent (dbremote).

#### Prerequisites

You must have the SYS\_RUN\_REPLICATION\_ROLE system role. The maximum message length, as defined by the -l option, must be the same for all databases in the system.

#### Procedure

1. Ensure that at least one remote user has neither a SEND AT nor a SEND EVERY option in their remote properties.

If *all* of your remote users have a SEND AT or a SEND EVERY clause defined, and you want to send and receive messages and then shut down, you must start the SQL Remote Message Agent (dbremote) using the -b option.

2. Start the SQL Remote Message Agent (dbremote).

On Microsoft Windows, the SQL Remote Message Agent (dbremote) is named `dbremote.exe`. On UNIX and Linux, the name is `dbremote`.

For example, the following command runs dbremote in batch mode on a database file named `c:\mydata.db`, connecting with user name `ManagerSteve` and password `passwd`:

```
dbremote -b -c "UID=ManagerSteve;PWD=passwd;DBF=c:\mydata.db"
```

The SQL Remote Message Agent (dbremote) receives and processes incoming messages, scans the transaction log once, creates and sends the outgoing messages, and then stops.

#### Results

The SQL Remote Message Agent (dbremote) is started.

### 1.4.3.4 Run the SQL Remote Message Agent (dbremote) on UNIX/Linux

On UNIX and Linux platforms, you run the SQL Remote Message Agent (dbremote) as a daemon by supplying the -ud option.

You must have the SYS\_RUN\_REPLICATION\_ROLE system role. The maximum message length, as defined by the -l option, must be the same on all databases in the system.

#### Related Information

[SQL Remote Message Agent \(dbremote\) Requirements \[page 97\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

### 1.4.4 SQL Remote Performance

SQL Remote divides the workload for notifications on subscriptions and publications between the database server and the SQL Remote Message Agent (dbremote) to manage network performance.

Each time a row in a table is inserted, deleted, or updated, a message is created for those users subscribed to the row. In addition, an update may cause the subscription expression to change, so that the statement is sent to some subscribers as a delete, some as an update, and some as an insert.

#### The Database Server

The database server handles publications.

#### The SQL Remote Message Agent (dbremote)

The **SQL Remote Message Agent (dbremote)** handles subscriptions.

The SQL Remote Message Agent (dbremote) reads the evaluated subscription expressions or subscription column entries from the transaction log, and matches the before and after values against the subscription value for each subscriber to the publication. In this way, the SQL Remote Message Agent (dbremote) sends the correct operations to each subscriber.

While a large number of subscribers do not have any effect on database server performance, they can affect SQL Remote Message Agent (dbremote) performance. The work in matching subscription values against large numbers of subscription values, and the work in sending the messages, can be demanding.

**In this section:**

[Tasks to Receive Messages \[page 103\]](#)

The SQL Remote Message Agent (dbremote) performs a series of tasks when receiving messages.

[Troubleshooting Performance When Receiving Messages \[page 104\]](#)

The major bottleneck for total throughput in a SQL Remote system is generally receiving messages from many remote databases and applying them to the database.

[Tasks to Send Messages \[page 110\]](#)

SQL Remote Message Agent (dbremote) executes a series of tasks when sending messages.

[Performance When Sending Messages \[page 111\]](#)

The major performance issue for sending messages is the turnaround time between when the data is entered at one site to when it appears at other sites.

## Related Information

[Transaction Log-based Replication \[page 38\]](#)

### 1.4.4.1 Tasks to Receive Messages

The SQL Remote Message Agent (dbremote) performs a series of tasks when receiving messages.

#### Polling for Incoming Messages

To check for new messages that have arrived at a database, the SQL Remote Message Agent (dbremote) polls for new messages.

#### Reading the messages

When messages arrive, they are read and stored in cache memory by the SQL Remote Message Agent (dbremote) until they can be applied.

If a message is missing and the SQL Remote Message Agent (dbremote) is running in continuous mode, then the SQL Remote Message Agent (dbremote) waits for the message to arrive in a subsequent poll. The number of polls that SQL Remote Message Agent (dbremote) waits is referred to as its **patience**, and is specified by the `-rp` option.

- If the missing message arrives before the SQL Remote Message Agent (dbremote) patience expires, then the missing message is added, in the correct order, to the cache.
- If the missing message does not arrive and the SQL Remote Message Agent (dbremote) patience expires, then the SQL Remote Message Agent (dbremote) sends a request to re-send the message from the publisher database.

Messages continue to be read and added to the cache until the cache memory usage is exceeded. When the cache memory usage specified using the `-m` option is exceeded, messages are deleted.

#### Applying the messages

The SQL Remote Message Agent (dbremote) applies the messages, in the correct order, to the subscriber database.

#### Waiting for confirmation that the messages are applied on the subscriber databases

Once the message has been received and applied at the subscribed database, confirmation is sent back to the publisher. When the publisher SQL Remote Message Agent (dbremote) receives the confirmation, it keeps track of the confirmation in a system table.

## Related Information

[Poll Interval Adjustments to Check for New Messages \[page 105\]](#)

[Throughput Adjustments by Caching Received Messages \[page 106\]](#)

[Request Adjustments to Resend Messages \[page 107\]](#)

[Database Worker Threads \[page 108\]](#)

[Guaranteed Message Delivery System \[page 114\]](#)

### 1.4.4.2 Troubleshooting Performance When Receiving Messages

The major bottleneck for total throughput in a SQL Remote system is generally receiving messages from many remote databases and applying them to the database.

To reduce this lag time, when running a SQL Remote Message Agent (dbremote) in continuous mode, you can adjust the following variables:

- How often the SQL Remote Message Agent (dbremote) checks for incoming messages.
- How much memory is used by the SQL Remote Message Agent (dbremote) for holding messages to be sent.
- How long the SQL Remote Message Agent (dbremote) waits for an out-of-order message to arrive before requesting that the message be re-sent.
- How many worker threads are used to process the received messages.

#### In this section:

[Poll Interval Adjustments to Check for New Messages \[page 105\]](#)

To check for new messages that have arrived at a database, the SQL Remote Message Agent (dbremote) polls for new messages. You can adjust the frequency of this polling.

[Throughput Adjustments by Caching Received Messages \[page 106\]](#)

When messages arrive, the SQL Remote Message Agent (dbremote) reads the messages and then stores them in cache memory until they are applied.

[Request Adjustments to Resend Messages \[page 107\]](#)

When a message is missing from a sequence, the SQL Remote Message Agent (dbremote) waits a specified number of polls before requesting that the missing message be re-sent.

[Database Worker Threads \[page 108\]](#)

There are a few points to consider when setting the number of database worker threads.



## 1.4.4.2.1 Poll Interval Adjustments to Check for New Messages

To check for new messages that have arrived at a database, the SQL Remote Message Agent (dbremote) polls for new messages. You can adjust the frequency of this polling.

The default polling interval from the end of one poll to the start of another is 1 minute. You can configure the polling interval using the `-rd` option, but the default is generally sufficient.

### Increasing the Polling Interval

You can poll more frequently by using a value in seconds. For example, the following command polls every thirty seconds:

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -rd 30s
```

In general, do not use a small polling interval unless you have a specific reason for requiring a very quick response time for messages. Setting a very small interval can have a detrimental effect on overall system throughput because:

- You can waste resources polling when no messages are in the queue. For example, if you are using email, each poll of the mail server places a load on your message system. Too frequent polling may affect your message system and produce no benefits.
- You can overload your system with resend requests. When adjusting the polling interval, also adjust the SQL Remote Message Agent (dbremote) patience. The patience is the number of polls the SQL Remote Message Agent (dbremote) waits for an out-of-sequence message to arrive before requesting that it be sent again.

### Decreasing the Polling Interval

You can poll less frequently, as in the following command, which polls every five minutes:

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -rd 5
```

Setting larger polling intervals can provide a better overall throughput of messages in your system, it can increase the time it takes to apply an individual messages. For example, if your polling period for incoming messages is too long, compared to the frequency with which messages are arriving, you can end up with messages sitting in the queue, waiting to be processed.

### Related Information

[Throughput Adjustments by Caching Received Messages \[page 106\]](#)

[Request Adjustments to Resend Messages \[page 107\]](#)

## 1.4.4.2.2 Throughput Adjustments by Caching Received Messages

When messages arrive, the SQL Remote Message Agent (dbremote) reads the messages and then stores them in cache memory until they are applied.

This caching of messages prevents:

- Rereading of out-of-order messages from the message system, which may lower performance on large systems. Caching of messages is useful when messages are being read over a WAN (such as Remote Access Services or POP3 through a modem).
- Contention between database worker threads reading messages (a single threaded task) because the message contents are cached.

### How Messages are Cached

Messages are stored in memory until they are applied by the SQL Remote Message Agent (dbremote) when one of the following conditions occurs:

- The transactions are so large that they require multi-part messages.
- The messages arrive out of order.

### Specifying the Message Cache Size

Use the SQL Remote Message Agent (dbremote) -m option to specify the size of the message cache. The -m option specifies the maximum amount of memory to be used by the SQL Remote Message Agent (dbremote) for holding messages. The allowed size can be specified as *n* (in bytes), *nK*, or *nM*. The default is 2048K (2M). When the specified cache memory usage is exceeded, messages are deleted.

The -m option is useful when you have a single consolidated database and a large number of remote databases.

### Example

The following command starts a SQL Remote Message Agent (dbremote) using 12 MB of memory as a message cache:

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -m 12M
```

## Related Information

[Poll Interval Adjustments to Check for New Messages \[page 105\]](#)

[Request Adjustments to Resend Messages \[page 107\]](#)

[Database Worker Threads \[page 108\]](#)

[Performance When Sending Messages \[page 111\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

### 1.4.4.2.3 Request Adjustments to Resend Messages

When a message is missing from a sequence, the SQL Remote Message Agent (dbremote) waits a specified number of polls before requesting that the missing message be re-sent.

The number of polls that the SQL Remote Message Agent (dbremote) waits is referred to as its patience. By default, the SQL Remote Message Agent (dbremote) has a patience of 1.

If the SQL Remote Message Agent (dbremote) has a patience of 1 and it expects to receive message 6 but it receives message 7, the SQL Remote Message Agent (dbremote) takes no action. Instead, the SQL Remote Message Agent (dbremote) waits for the results of the next poll. If after the next poll, Message 6 is still missing, then the SQL Remote Message Agent (dbremote) issues a resend request for Message 6.

### Increasing the Resend Patience

Suppose you have a very small polling interval, and a message system that does not preserve the order in which messages arrive. It may be common for out-of-sequence messages to arrive after two or three polls have been completed. In this example, use the `-rp` option to increase the SQL Remote Message Agent (dbremote) patience so that a large number of unnecessary resend requests are not sent. The `-rp` option is often used with the `-rd` option that sets the polling interval.

### Example

There are two remote users, named `user1` and `user2`, both of which run the SQL Remote Message Agent (dbremote) with a polling interval of 30 seconds and a patience of 3 polls. For example, they use the following command to run their SQL Remote Message Agents (dbremote):

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -rd 30s -rp 3
```

In the following sequence of operations, messages are marked as `userX.n` where `X` is the user name and `n` is the message number. For example, `user1.5` is the fifth message from `user1`. The SQL Remote Message Agent (dbremote) expects messages to start at number 1 for both users.

At time 0 seconds:

1. The SQL Remote Message Agent (dbremote) reads user1.1, user2.4
2. The SQL Remote Message Agent (dbremote) applies user1.1
3. The SQL Remote Message Agent (dbremote) patience is now user1: N/A, user2: 3, as an out of sequence message has arrived from user 2

At time 30 seconds:

1. The SQL Remote Message Agent (dbremote) reads: no new messages
2. The SQL Remote Message Agent (dbremote) applies: nothing
3. The SQL Remote Message Agent (dbremote) patience is now user1: N/A, user2: 2

At time 60 seconds:

1. The SQL Remote Message Agent (dbremote) reads: user1.3
2. The SQL Remote Message Agent (dbremote) applies: no new messages
3. The SQL Remote Message Agent (dbremote) patience: user1: 3, user2: 1

At time 90 seconds:

1. The SQL Remote Message Agent (dbremote) reads: user1.4
2. The SQL Remote Message Agent (dbremote) applies: none
3. The SQL Remote Message Agent (dbremote) patience user1: 3, user2: 0
4. The SQL Remote Message Agent (dbremote) issues resend to user2

When a user receives a new message, it resets the SQL Remote Message Agent (dbremote) patience even if that message is not the one expected.

At time 120 seconds:

1. The SQL Remote Message Agent (dbremote) reads: user1.2 and user2.2
2. The SQL Remote Message Agent (dbremote) applies user1.2, user1.3, user1.4, and user2.2
3. The SQL Remote Message Agent (dbremote) patience user1: N/A, user2: N/A

## Related Information

[Poll Interval Adjustments to Check for New Messages \[page 105\]](#)

[Throughput Adjustments by Caching Received Messages \[page 106\]](#)

[Database Worker Threads \[page 108\]](#)

[Performance When Sending Messages \[page 111\]](#)

### 1.4.4.2.4 Database Worker Threads

There are a few points to consider when setting the number of database worker threads.

1. It reads the messages. Messages are read and the header information is examined (to determine the correct order of application). Reading messages from the message system is single-threaded.
2. It applies the messages. Read messages are passed off to database worker threads to be applied.

On remote databases, the messages are usually applied serially. In a multi-tier system, a remote database can also be a consolidated database for other remotes. On this type of a remote database, the messages are applied as on a consolidated database.

On the consolidated database, the default is to apply the messages serially. You can use additional database worker threads to apply incoming messages from remote users in parallel using the `-w` option of the SQL Remote Message Agent utility (`dbremote`).

When database worker threads are used on a consolidated database:

- Messages from different remote users are applied in parallel.
- Messages from a single remote user are applied serially.  
For example, ten messages from a single remote user are applied by a single worker thread in the correct order.

## Advantages of Using Database Worker Threads

Using database worker threads on the consolidated database can improve throughput by allowing messages to be applied in parallel rather than serially. The performance advantage is significant when the database server is on a system with a striped drive array.

## Disadvantages of Using Database Worker Threads

Using database worker threads on the consolidated database can decrease throughput if they cause a lot of locking between users.

A deadlock is handled by re-applying the rolled-back transaction later.

## Setting the Number of Database Worker Threads

On the consolidated database, use the `-w` option to set the number of database worker threads. For example, the following command sets the number of worker threads to 5:

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -w 5
```

## Related Information

[Poll Interval Adjustments to Check for New Messages \[page 105\]](#)

[Throughput Adjustments by Caching Received Messages \[page 106\]](#)

[Request Adjustments to Resend Messages \[page 107\]](#)

[Running the SQL Remote Message Agent in Continuous Mode \(`dbremote`\) \[page 97\]](#)

[SQL Remote Message Agent Utility \(`dbremote`\) \[page 215\]](#)

[Performance When Sending Messages \[page 111\]](#)

## 1.4.4.3 Tasks to Send Messages

SQL Remote Message Agent (dbremote) executes a series of tasks when sending messages.

### Scanning the publisher transaction log

The SQL Remote Message Agent (dbremote) scans the transaction log of the publisher database and translates the transaction log entries into messages for subscribers. The maximum message length, as defined by the -l option, must be the same on all databases in the system.

For large transactions, the SQL Remote Message Agent (dbremote) creates multi-part messages. These messages each contain a sequence number that keeps track of their place in the transaction. The SQL Remote Message Agent (dbremote) on the subscriber database uses the sequence number to ensure that the messages are applied in the correct order and that no message is lost.

### Sending messages to the remote databases

The SQL Remote Message Agent (dbremote) sends messages at times specified by the send frequency properties of each remote user.

The SQL Remote Message Agent (dbremote) sends messages earlier if its cache memory exceeds the set value. The SQL Remote Message Agent (dbremote) stores its messages in cache memory. When the cache memory being used exceeds the specified value, messages are sent.

### Processing resend requests from remote databases

When a user requests that a message be re-sent, the SQL Remote Message Agent (dbremote) on the publisher database interrupts the regular message sending process to process the resend request.

You control the urgency with which these resend requests are processed with the -ru option.

### Sending confirmations to the publisher database

Once a message has been received and applied at the subscribed database, confirmation is sent back to the publisher.

## Related Information

[Send Delay Adjustments \[page 111\]](#)

[Throughput Adjustments by Caching Sent Messages \[page 112\]](#)

[Re-send Request Processing Speed \[page 113\]](#)

[Guaranteed Message Delivery System \[page 114\]](#)

## 1.4.4.4 Performance When Sending Messages

The major performance issue for sending messages is the turnaround time between when the data is entered at one site to when it appears at other sites.

To reduce this lag time, when sending messages with the SQL Remote Message Agent (dbremote), you can adjust the following variables:

- How often messages are sent to remote databases.
- The size of the messages.
- How quickly resend requests are processed.

### In this section:

#### [Send Delay Adjustments \[page 111\]](#)

To create messages to send, the SQL Remote Message Agent (dbremote) polls for new data from the transaction log. The send delay is the time to wait between polls for more transaction log data to send.

#### [Throughput Adjustments by Caching Sent Messages \[page 112\]](#)

The SQL Remote Message Agent (dbremote) caches messages to be sent in a configurable area of memory.

#### [Re-send Request Processing Speed \[page 113\]](#)

Because resending a message interrupts the regular message sending process, the SQL Remote Message Agent (dbremote) delays processing resend requests.

## Related Information

[Troubleshooting Performance When Receiving Messages \[page 104\]](#)

### 1.4.4.4.1 Send Delay Adjustments

To create messages to send, the SQL Remote Message Agent (dbremote) polls for new data from the transaction log. The send delay is the time to wait between polls for more transaction log data to send.

The default polling interval from the end of one poll to the start of another is 1 minute. You can configure the send delay using the `-sd` option, but the default is generally sufficient. The send delay should be less than or equal to the remote users' send frequency.

## Decreasing the Send Delay

You can poll more frequently by using a value in seconds. For example, the following command polls every thirty seconds:

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -sd 30s ...
```

## Increasing the Send Delay

You can poll less frequently, as in the following command, which polls every 60 minutes:

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -sd 60
```

Typically, larger send intervals mean that the SQL Remote Message Agent (dbremote) does most of the message creation work before sending the messages. Smaller intervals are generally preferred as they spread out the message creation work.

## Related Information

[Throughput Adjustments by Caching Sent Messages \[page 112\]](#)

[Re-send Request Processing Speed \[page 113\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

### 1.4.4.4.2 Throughput Adjustments by Caching Sent Messages

The SQL Remote Message Agent (dbremote) caches messages to be sent in a configurable area of memory.

When all remote databases are receiving unique subsets of the operations being replicated, a separate message for each remote database is built up concurrently. Only one message is built for a group of remote users that is receiving the same operations. Messages are sent when:

- The send frequency is reached.
- When the cache memory being used exceeds the -m value.
- When the size of the message reaches its maximum size (as specified by the -l option).

## Specifying the message cache size

The size of the message cache is specified on the SQL Remote Message Agent (dbremote) command, using the -m option.

The -m option specifies the maximum amount of memory to be used by the SQL Remote Message Agent (dbremote) for building messages. The allowed size can be specified as *n* (in bytes), *nK*, or *nM*. The default is 2048K (2M).

The -m option is useful when you have a single consolidated database and a large number of remote databases.



## Example

The following command starts a SQL Remote Message Agent (dbremote) using 12 MB of memory as a message cache:

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -m 12M
```

## Related Information

[Send Delay Adjustments \[page 111\]](#)

[Re-send Request Processing Speed \[page 113\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

### 1.4.4.4.3 Re-send Request Processing Speed

Because resending a message interrupts the regular message sending process, the SQL Remote Message Agent (dbremote) delays processing resend requests.

By default, the SQL Remote Message Agent (dbremote) waits for a time that is half of the send frequency of the remote user who requested the resend.

To resend a message, the SQL Remote Message Agent (dbremote) does the following tasks:

- It stops scanning the transaction log and stops building new messages.
- It deletes the current messages that are stored in its cache waiting to be sent. All of the work that the SQL Remote Message Agent (dbremote) did in reading the transaction log and building those messages is lost.
- It re-reads the transaction log from the offset requested in the resend request. The SQL Remote Message Agent (dbremote) builds the messages and stores them in its cache.
- It waits until the next send frequency time occurs and then sends the messages.

You must balance the urgency of sending requests for re-sent messages with the priority of processing regular messages.

The `-ru` option controls the urgency of the resend requests. To delay processing resend requests until more have arrived, set this option to a longer time. For example, the following command waits one hour before processing a resend request:

```
dbremote -c "DSN=SQL Anywhere 17 Demo;PWD=sql" -ru 1h
```

## Related Information

[Send Delay Adjustments \[page 111\]](#)

[Throughput Adjustments by Caching Sent Messages \[page 112\]](#)

## 1.4.5 Guaranteed Message Delivery System

The Guaranteed Message Delivery System ensures that all replicated operations are applied in the correct order, no replicated operations are missed, and no replicated operation is applied twice.

The Guaranteed Message Delivery System uses the following information:

### The status information maintained in the SYSREMOTUSER system table

This table contains a row for each subscriber, with status information for messages sent to and received by that subscriber. For example:

- On the consolidated database, the SYSREMOTUSER system table contains a row for each remote user.
- On each remote database, the SYSREMOTUSER system table contains a single row with information about the consolidated database.

The SYSREMOTUSER system table is maintained by the SQL Remote Message Agent (dbremote).

On the subscriber database, the SQL Remote Message Agent (dbremote) sends a confirmation to the publisher database to ensure that the SYSREMOTUSER system table is maintained correctly at each end of the subscription.

### The information in the header of the messages

The SQL Remote Message Agent (dbremote) reads the header information in the messages and uses this information to update the SYSREMOTUSER system table. A message includes the following information in its header:

#### Its resend\_count

A counter that keeps track of the number of times the receiving database has lost messages.

In the following example, the resend\_count is 1.

```
Current message's header: (1-0000942712-0001119170-0)
```

#### The transaction log offset of the last COMMIT in the previous message

In the following example, the transaction log offset of the last commit in the previous message is 0000942712.

```
Previous message's header: (0-0000923357-0000942712-0)  
Current message's header: (0-0000942712-0001119170-0)
```

#### The transaction log offset of the last COMMIT in the current message

In the following example, the last commit in the current message is 0001119170:

```
Current message's header: (0-0000942712-0001119170-0)
```

If a transaction spans several messages, both transaction log offsets can be identical until the final message contains a COMMIT.

In the following example, the COMMIT does not occur until the fourth message:

```
(0-0000942712-0000942712-0)
(0-0000942712-0000942712-1)
(0-0000942712-0000942712-2)
(0-0000942712-0001119170-3)
```

### A sequence number

When a transaction spans several messages, this sequence number is used to order the messages correctly.

A sequence number of zero can indicate that:

- The message is not part of a multi-part message if the transaction log offsets are different. In the following example, the messages are not part of a multi-part message:

```
(0-0000923200-0000923357-0)
(0-0000923357-0000942712-0)
```

- The message is the first of a multi-part message if the transaction log offsets are the same. In the following example, the first message is part of a multi-part message:

```
(0-0000942712-0000942712-0)
(0-0000942712-0000942712-1)
(0-0000942712-0000942712-2)
(0-0000942712-0001119170-3)
```

### In this section:

#### [Order of Operations \[page 115\]](#)

To ensure that the replicated statements are applied in the correct order, the Guaranteed Message Delivery System uses the transaction log offsets of the publisher and subscriber databases.

#### [Lost or Corrupt Messages \[page 117\]](#)

The SYSREMOTEUSER system table provides for the management of resending messages.

#### [Messages Are Applied Only Once \[page 118\]](#)

To ensure that operations are not applied more than once, the subscriber SQL Remote Message Agent (dbremote) compares the resend\_count value in a messages header with the rereceive\_count in its local SYSREMOTEUSER system table. If the resend\_count value is smaller than rereceive\_count, the message is not applied; it is deleted.

## 1.4.5.1 Order of Operations

To ensure that the replicated statements are applied in the correct order, the Guaranteed Message Delivery System uses the transaction log offsets of the publisher and subscriber databases.

Each COMMIT is marked in the transaction log by a well-defined offset. The order of transactions can be determined by comparing their transaction log offset values. Each message includes the following transaction log offsets:

- The transaction log offset of the last COMMIT in the previous message. If a transaction spans several messages, there is a sequence number within the transaction to order the messages correctly.

- The transaction log offset of the last COMMIT in the message.

## Message Ordering

When messages are sent, they are ordered by the offset of the last COMMIT of the preceding message. If a transaction spans several messages, a sequence number within the transaction is used to order the messages correctly.

## Sending Messages

The `log_sent` column in the `SYSREMOTUSER` system table holds the local transaction log offset for the last message sent to the subscriber.

The following describes how the `SYSREMOTUSER` system tables are updated when messages are sent.

1. When the publisher SQL Remote Message Agent (`dbremote`) sends a message to a subscriber, it also sets the `log_sent` value to the transaction log offset value of the last COMMIT in the sent message. For example, the publisher sends the following message to `user1`.

```
(0-0000923200-0000923357-0)
```

In the publisher's `SYSREMOTUSER` system table, the publisher sets the `log_sent` value to 0000923357 for `user1`.

2. When the message is received and applied at the subscriber database, a confirmation is sent to the publisher. The confirmation includes the last transaction log offset that was applied by the subscriber database. For example, the message confirms that `user1` applied all of the transactions up to and including the transaction log offset 0000923357.
3. When the publisher SQL Remote Message Agent (`dbremote`) receives the confirmation, it sets the `confirm_sent` column to the value of the confirmation offset for the user in the `SYSREMOTUSER` system table. For example, the publisher sets the `confirm_sent` column to 0000923357 for `user1` in the publisher's `SYSREMOTUSER` system table.

Both the `log_sent` and `confirm_sent` values contain transaction log offsets of the publisher's transaction log. The `confirm_sent` value cannot be a later offset than `log_sent` value.

## Receiving Messages

The following describes how the `SYSREMOTUSER` system tables are updated when messages are received.

1. When the SQL Remote Message Agent (`dbremote`) at a subscriber database receives and applies a replication update, it updates the `log_received` column in the `SYSREMOTUSER` system table with the offset of the last COMMIT in the message.

For example, when the subscriber receives and applies the following message, the `log_received` value in the `SYSREMOTEUSER` system table is set to 0000923357.

```
(0-0000923200-0000923357-0)
```

The `log_received` column at any subscriber database contains a transaction log offset that exists in the publisher database transaction log.

2. When the operations are received and applied, the subscriber SQL Remote Message Agent (`dbremote`) sets the `confirm_received` value in its `SYSREMOTEUSER` system table, and then sends confirmation to the publisher database.

## Related Information

[Lost or Corrupt Messages \[page 117\]](#)

[Messages Are Applied Only Once \[page 118\]](#)

### 1.4.5.2 Lost or Corrupt Messages

The `SYSREMOTEUSER` system table provides for the management of resending messages.

The `SYSREMOTEUSER` system table contains two columns to manage resending messages:

#### **resend\_count** column

A counter that keeps track of the number of times that the subscriber database has lost messages.

#### **rereceive\_count** column

A counter that keeps track of the number of times the SQL Remote Message Agent (`dbremote`) has determined that messages from a publisher user have been lost.

When messages are received in the proper order at a subscriber database:

1. The subscriber SQL Remote Message Agent (`dbremote`) applies the messages in the correct order and updates its `SYSREMOTEUSER` system table.
2. The subscriber SQL Remote Message Agent (`dbremote`) sends a confirmation message to the publisher.
3. When the publisher receives the confirmation, its SQL Remote Message Agent (`dbremote`) updates its `SYSREMOTEUSER` system table.

When messages are not received in the proper order:

1. The subscriber SQL Remote Message Agent (`dbremote`) sends a resend request and increments the `rereceive_count` value in its `SYSREMOTEUSER` system table.
2. When the publisher receives the resend request, it increments the `resend_count` value in its `SYSREMOTEUSER` system table for the subscriber.
3. In the publisher's `SYSREMOTEUSER` system table, the `log_sent` value is set to the value in the `confirm_sent` column. Resetting of the `log_sent` value causes operations to be re-sent.

## Related Information

[Order of Operations \[page 115\]](#)

[Messages Are Applied Only Once \[page 118\]](#)

### 1.4.5.3 Messages Are Applied Only Once

To ensure that operations are not applied more than once, the subscriber SQL Remote Message Agent (dbremote) compares the resend\_count value in a messages header with the rereceive\_count in its local SYSREMOTEEUSER system table. If the resend\_count value is smaller than rereceive\_count, the message is not applied; it is deleted.

## Related Information

[Order of Operations \[page 115\]](#)

[Lost or Corrupt Messages \[page 117\]](#)

### 1.4.6 Message Size

The SQL Remote Message Agent (dbremote) provides a number of encoding and compression features.

#### Compatibility

The system can be set to be compatible with previous versions of SQL Anywhere.

#### Compression

You can select a level of compression for your messages.

Message size affects the efficiency with which messages pass through a system. Compressed messages can be processed more efficiently by a message system than uncompressed messages. However, compression can take a significant amount of time.

#### Encoding

SQL Remote encodes messages to ensure that they pass through message systems uncorrupted. The encoding scheme can be customized to provide extra features.

#### In this section:

[Prevention of Message Corruption with Encoding \[page 119\]](#)

SQL Remote encodes messages to ensure that they pass through message systems uncorrupted.

## Related Information

[SQL Remote Upgrades](#)

[SQL Remote Message Agent \(dbremote\) Requirements \[page 97\]  
compression Option \[SQL Remote\]](#)

### 1.4.6.1 Prevention of Message Corruption with Encoding

SQL Remote encodes messages to ensure that they pass through message systems uncorrupted.

The default message-encoding behavior of SQL Remote is as follows:

- If the message system can use binary message formats, the messages are not encoded.
- If the message system, for example SMTP, requires text-based message formats, then an encoding DLL (`dbencod17.dll`) translates the messages into a text format before sending. The message format is unencoded at the receiving end using the same DLL. You can customize the encoding scheme to provide extra features.
- If the compression database option is set to -1, then a version 5 compatible encoding is used for all message systems.

#### In this section:

[Custom Encoding Schemes \[page 119\]](#)

To implement a custom encoding scheme, you can build a custom encoding DLL. You can use this custom DLL to apply special features required for a particular messages system, or to collect statistics, such as how many messages are sent to each user.

## Related Information

[SQL Remote Upgrades](#)

### 1.4.6.1.1 Custom Encoding Schemes

To implement a custom encoding scheme, you can build a custom encoding DLL. You can use this custom DLL to apply special features required for a particular messages system, or to collect statistics, such as how many messages are sent to each user.

The header file `%SQLANY17%\SDK\Include\dbrmt.h` includes an application programming interface that you can use to build a custom encoding scheme.

To use your custom DLL, set the message control parameter `encode_dll` to a value that is the full path to the custom DLL. For example:

```
SET REMOTE FTP OPTION "Public"."encode_dll" = 'c:\sqlany17\Bin32\custom.dll';
```

## i Note

Encoding and decoding must be compatible: if you implement a custom encoding, you must make sure that the DLL is present at the receiving end, and that the DLL is in place to decode your messages properly.

## Related Information

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

## 1.4.7 SQL Remote Message Systems

In SQL Remote replication, a message system is a protocol for exchanging messages between the consolidated database and a remote database. SQL Remote exchanges data among databases using one or more underlying message systems.

SQL Remote supports the following message systems:

### **File sharing**

A simple system requiring no extra software.

### **FTP**

Internet file transfer protocol.

### **HTTP**

Hypertext transfer protocol.

### **SMTP/POP**

Internet email protocol.

You choose a message system when you assign REMOTE or CONSOLIDATE privilege to a user.

Each message system that is used in a SQL Remote system has control parameters and other settings that must be set up.

Not all message systems are supported on all operating systems.

### **In this section:**

#### [Message System Setup \[page 121\]](#)

Before you can use a message system, you must set the publisher's address.

#### [Creating a Message Type \(SQL Central\) \[page 122\]](#)

Add message types for SQL Remote users.

#### [Altering a Message Type \(SQL Central\) \[page 123\]](#)

Alter message types to change the address of a publisher. You cannot change the name of an existing message type; instead, you must delete it and create a new message type with the new name.

#### [Deleting a Message Type \(SQL Central\) \[page 123\]](#)

Delete message types, which removes the publisher address from the message definition.



#### [Remote Message Type Control Parameters \[page 124\]](#)

Message control parameters are held in the database. To set the control parameter, execute a SET REMOTE OPTION statement.

#### [The FILE Message System \[page 125\]](#)

SQL Remote can be used even if you do not have an email or FTP system in place, by using the FILE message system.

#### [The FTP Message System \[page 127\]](#)

In the FTP message system, messages are stored in directories under the root directory of an FTP host. The host and root directory are specified by message system control parameters in the registry or initialization file, and the address of each user is the subdirectory where their messages are held.

#### [The HTTP Message System \[page 130\]](#)

Using the HTTP message system, SQL Remote sends messages over the Internet using the Hypertext Transfer Protocol (HTTP). The messages are encoded in a text format and sent via HTTP to the target database.

#### [The SMTP Message System \[page 134\]](#)

With the SMTP system, SQL Remote sends messages using Internet mail.

## Related Information

[Supported Platforms](#)

[Granting REMOTE Privilege \(SQL Central\) \[page 30\]](#)

[Granting CONSOLIDATE Privilege \(SQL Central\) \[page 33\]](#)

### 1.4.7.1 Message System Setup

Before you can use a message system, you must set the publisher's address.

Each message type definition includes the message system type name (FILE, FTP, HTTP, or SMTP) and the address of the publisher under that message type.

The address supplied with a message type definition is closely tied to the publisher ID of the database.

### Extraction uUtility (dbxtract)

The publisher address at a consolidated database is used by the Extraction utility (dbxtract) and the *Extract Database Wizard* as a return address when creating remote databases. It is also used by the SQL Remote Message Agent (dbremote) to identify the location of incoming messages for the FILE system.

## Related Information

[The FILE Message System \[page 125\]](#)

[The FTP Message System \[page 127\]](#)

[The SMTP Message System \[page 134\]](#)

[The HTTP Message System \[page 130\]](#)

## 1.4.7.2 Creating a Message Type (SQL Central)

Add message types for SQL Remote users.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click *SQL Remote Users*.
3. Click the *Message Types* tab.
4. Click **File** > **New** > **Message Type**.
5. Type a name for the message type. The name should correspond to a message-type DLL already installed in your SQL Anywhere installation directory. Click *Next*.
6. In the *What Is The Publisher Address* field, type a publisher address. Click *Finish*.

### Results

The new message type is created.

## Related Information

[Altering a Message Type \(SQL Central\) \[page 123\]](#)

[Deleting a Message Type \(SQL Central\) \[page 123\]](#)

[CREATE REMOTE \[MESSAGE\] Statement \[SQL Remote\]](#)

### 1.4.7.3 Altering a Message Type (SQL Central)

Alter message types to change the address of a publisher. You cannot change the name of an existing message type; instead, you must delete it and create a new message type with the new name.

#### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

#### Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click *SQL Remote Users*.
3. Click the *Message Types* tab.
4. Right-click the message type you want to alter and click *Properties*.
5. Update the message type properties and click *OK*.

#### Results

The message properties are updated.

#### Related Information

[Creating a Message Type \(SQL Central\) \[page 122\]](#)

[Deleting a Message Type \(SQL Central\) \[page 123\]](#)

[ALTER REMOTE \[ MESSAGE \] Statement \[SQL Remote\]](#)

### 1.4.7.4 Deleting a Message Type (SQL Central)

Delete message types, which removes the publisher address from the message definition.

#### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Procedure

1. In SQL Central, use the [SQL Anywhere 17](#) plug-in to connect to the database.
2. Double-click [SQL Remote Users](#).
3. Click the [Message Types](#) tab.
4. Right-click the message type you want to remove and click [Delete](#).
5. Click [Yes](#).

## Results

The message type is deleted.

## Related Information

[Creating a Message Type \(SQL Central\) \[page 122\]](#)

[Altering a Message Type \(SQL Central\) \[page 123\]](#)

[DROP REMOTE \[MESSAGE\] Statement \[SQL Remote\]](#)

### 1.4.7.5 Remote Message Type Control Parameters

Message control parameters are held in the database. To set the control parameter, execute a SET REMOTE OPTION statement.

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

For example, the following statement sets the FTP host to `ftp.mycompany.com` for the FTP link for user `myuser`:

```
SET REMOTE FTP OPTION myuser.host = 'ftp.mycompany.com';
```

To view the message link parameters (using SQL), query the SYSREMOTEOPTION system view:

```
SELECT * from SYSREMOTEOPTION;
```

### Message Link Parameters Stored on Disk

Earlier versions of SQL Remote stored the message link parameters outside the database. The external storage of message link parameters is not recommended.

The message link control parameters can be stored in the following locations:

#### Windows only

In the registry, at one of the following locations:

```
HKEY_CURRENT_USER\Software\SAP\SQL Remote
```

or

```
HKEY_LOCAL_MACHINE\Software\SAP\SQL Remote
```

The parameters for each message link are stored in a key under the SQL Remote key, with the name of the message link. Here are some examples.

```
HKEY_CURRENT_USER\Software\SAP\SQL Remote\FILE\max_retries  
HKEY_LOCAL_MACHINE\Software\SAP\SQL Remote\SMTP\smtp_userid
```

#### Windows or Unix

You can use the SQLREMOTE environment variable to specify the FILE message system Directory option setting.

When the SQL Remote Message Agent (dbremote) loads a message link, the link uses the setting of the current publisher or, if a setting is not specified, of groups to which the publisher belongs.

On Windows, when the SQL Remote Message Agent (dbremote) that supports storing the message link parameters in the database is run for the first time, it copies the link options from the registry to the database.

## Related Information

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

[SYSREMOPTOPTION System View](#)

## 1.4.7.6 The FILE Message System

SQL Remote can be used even if you do not have an email or FTP system in place, by using the FILE message system.

### Addresses in the FILE Message System

The FILE message system is a simple FILE-sharing system. A FILE address for a remote user is a subdirectory into which all their messages are written. To retrieve messages, an application reads the messages from the directory containing the user's files. Return messages are sent to the address (written to the directory) of the consolidated database.

When the SQL Remote Message Agent (dbremote) is running as a service, the account it is running under must have permissions to read from and write to all necessary directories. If the correct permissions are not assigned, the SQL Remote Message Agent is unable to access network drives.

## Root Directory for Addresses

Typically, the FILE message system addresses are subdirectories of a shared directory that is available to all SQL Remote users, whether by modem or a local area network. Each user should have a registry entry, initialization file entry, or SQLREMOTE environment variable pointing to the shared directory.

You can also use the FILE system to put the messages in directories on the consolidated and remote computers. You can use a simple file transfer mechanism to exchange files and complete replication.

## FILE Message Control Parameters

The FILE message system uses the following control parameters that are set by the SET REMOTE OPTION statement:

### **directory**

The directory under which the messages are stored. This parameter is an alternative to the SQLREMOTE environment variable.

### **debug**

The setting for this parameter is either YES or NO. The default is NO. When set to YES, all FILE system calls made by the FILE link are displayed in the output log.

### **encode\_dll**

If you are using a custom encoding scheme, you must set this parameter to the full path of the custom encoding DLL that you created.

### **invalid\_extensions**

A comma-separated list of file extensions that you do not want the SQL Remote Message Agent (dbremote) to use when generating files in the messaging system.

### **max\_retries**

By default, when SQL Remote is running in continuous mode and an error occurs when accessing the message system, it shuts down after the send and/or receive phases. Use this parameter to specify the number of times you want SQL Remote to retry the send and/or receive phases before it shuts down.

### **pause\_after\_failure**

This parameter applies when the max\_retries parameter is specified to a value other than zero and SQL Remote is running in continuous mode. When an error occurs in the message system, this parameter defines the number of seconds SQL Remote waits between retrying the send and/or receive phases.

### **unlink\_delay**

The number of seconds to wait before attempting to delete a file if the previous attempt to delete the file failed. If no value is defined for unlink\_delay, then the default behavior is set to pause for 1 second after the first failed attempt, 2 seconds after the second failed attempt, 3 seconds after the third failed attempt, and 4 seconds after the fourth failed attempt.

### **output\_log\_send\_on\_error**

Sends log information when an error occurs.

### **output\_log\_send\_limit**

Limits the amount of information that is sent to the consolidated database. The `output_log_send_limit` option specifies the number of bytes at the end of the output log (that is, the most recent entries) that are sent to the consolidated database. The default is 5 KB.

### **output\_log\_send\_now**

When set to YES, sends output log information to the consolidated database. On the next poll, the remote database sends the output log information and then resets the `output_log_send_now` option to NO.

## **Related Information**

[The FTP Message System \[page 127\]](#)

[The SMTP Message System \[page 134\]](#)

[The HTTP Message System \[page 130\]](#)

[Message Size \[page 118\]](#)

[CREATE REMOTE \[MESSAGE\] Statement \[SQL Remote\]](#)

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

[SQLREMOTE Environment Variable](#)

## **1.4.7.7 The FTP Message System**

In the FTP message system, messages are stored in directories under the root directory of an FTP host. The host and root directory are specified by message system control parameters in the registry or initialization file, and the address of each user is the subdirectory where their messages are held.

### **FTP Message Control Parameters**

The FTP message system uses the following control parameters that are set by the SET REMOTE OPTION statement:

#### **host**

The host name of the computer where the FTP server is running. This parameter can be a host name (such as ftp.sap.com) or an IP address (such as 192.138.151.66).

#### **user**

The user name for accessing the FTP host.

#### **password**

The password for accessing the FTP host.

#### **root\_directory**

The root directory within the FTP host site that the messages are stored under.

**port**

The IP port number used for the FTP connection. This parameter is usually not required.

**debug**

This parameter is set either to YES or NO. The default is NO. When set to YES, debugging output is displayed in the output log.

**active\_mode**

This parameter controls how SQL Remote establishes the server/client connection. This parameter is set either to YES or NO. The default is NO (passive mode). Passive mode is the preferred transfer mode and the default for the FTP message link. In passive mode, all data transfer connections are initiated by the client, in this case, the message link. In active mode, the FTP server initiates all data connections.

**reconnect\_retries**

The number of times the link should try to open a socket with the server before failing. The default value is 4. When you set this parameter, only reconnections are affected. The initial connection made by the FTP link is not affected.

**reconnect\_pause**

The time in seconds to pause between each connection attempt. The default setting is 30 seconds. When you set this parameter, only reconnections are affected. The initial connection made by the FTP link is not affected.

**suppress\_dialogs**

This parameter is set to ON or OFF. The default value is OFF. If it is set to ON, the *Connect* window does not appear after failed attempts to connect to the FTP server. Instead, an error is generated.

**invalid\_extensions**

A comma-separated list of file extensions that you do not want dbremote to use when generating files in the messaging system.

**encode\_dll**

If you have implemented a custom encoding scheme, you must set this parameter to the full path of the custom encoding DLL that you created.

**max\_retries**

By default, when SQL Remote is running in continuous mode and an error occurs when accessing the message system, it shuts down after the send and/or receive phases. Use this parameter to specify the number of times you want SQL Remote to retry the send and/or receive phases before it shuts down.

**pause\_after\_failure**

This parameter applies when the max\_retries parameter is specified to a value other than zero and SQL Remote is running in continuous mode. When an error occurs in the message system, this parameter defines the number of seconds SQL Remote waits between retrying the send and/or receive phases.

**output\_log\_send\_on\_error**

Sends log information when an error occurs.

**output\_log\_send\_limit**

Limits the amount of information that is sent to the consolidated database. The output\_log\_send\_limit option specifies the number of bytes at the end of the output log (that is, the most recent entries) that are sent to the consolidated database. The default is 5 KB.



## **output\_log\_send\_now**

When set to YES, sends output log information to the consolidated database. On the next poll, the remote database sends the output log information and then resets the output\_log\_send\_now option to NO.

### **In this section:**

[Troubleshooting FTP Problems \[page 129\]](#)

Most problems with the FTP message link are caused by network system issues.

## **Related Information**

[The HTTP Message System \[page 130\]](#)

[Supported Platforms](#)

[Message Size \[page 118\]](#)

[The FILE Message System \[page 125\]](#)

[The SMTP Message System \[page 134\]](#)

[CREATE REMOTE \[MESSAGE\] Statement \[SQL Remote\]](#)

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

## **1.4.7.7.1 Troubleshooting FTP Problems**

Most problems with the FTP message link are caused by network system issues.

The following list describes tests you can use to troubleshoot problems.

### **Set the debug message control parameter**

Review the debug output to determine whether you are connecting to the FTP server. If you are connecting, the debug output should indicate which FTP commands are failing.

### **Ping the FTP server**

If the FTP link is not able to connect to the FTP server, test your system network configuration. For example, run the following command:

```
ping ftp-server-name
```

The IP address of the FTP server and the ping (round trip) time to the FTP server should be returned. If you cannot ping the FTP server, then you have a network configuration problem, and you should contact your network administrator.

### **Check that passive mode works**

If the FTP link is connecting to the FTP server, but is unable to open a data connection, make sure that an FTP client can use passive mode to transfer data with the server.

Passive mode is the preferred transfer mode and the default for the FTP message link. In passive mode, all data transfer connections are initiated by the client, in this case, the message link. In active mode, the FTP server initiates all data connections. If your FTP server is behind an incorrectly configured firewall, you may

not be able to use the default passive transfer mode because the firewall blocks socket connections to the FTP server on ports other than the FTP control port.

Using an FTP user program that allows you to set the transfer mode between **active** and **passive**, set the transfer mode to passive and try to upload or download a file. If the client you are using cannot transfer the file without using active mode then you should reconfigure the firewall and FTP server to allow passive mode transfers or set the `active_mode` message control parameter to YES. Active mode transfers may not work in all network configurations. For example, if your client is behind an IP masquerading gateway incoming connections may fail depending on your gateway software.

#### **Check permissions and directory structures**

If the FTP server is connecting and having problems getting directory listings or manipulating files, make sure your permissions are set up correctly and the required directories exist.

Log in to the FTP server using an FTP program. Change directories to the location stored in the `root_directory` parameter. If the directories you need do not appear, the `root_directory` control parameter may be set incorrectly or the directories may not exist.

Test permissions by fetching a file in your message directory and uploading a file to the consolidated database directory. If errors are returned, your FTP server permissions are set up incorrectly.

## **1.4.7.8 The HTTP Message System**

Using the HTTP message system, SQL Remote sends messages over the Internet using the Hypertext Transfer Protocol (HTTP). The messages are encoded in a text format and sent via HTTP to the target database.

Messages are sent and received using a SQL Anywhere database acting as an HTTP server.

### **Setting up the HTTP Message Server**

You use the SQL Anywhere database server to act as the HTTP server that transfers SQL Remote messages to and from remote databases. By default, a newly initialized SQL Anywhere database does not have the web services defined to allow it to act as a message server. Three system stored procedures, `sr_add_message_server`, `sr_drop_message_server` and `sr_update_message_server` are defined in newly created SQL Anywhere databases to allow you to define the required database objects so that the database can act as an HTTP Server to transfer SQL Remote messages.

#### **i Note**

The database must have been initialized with SQL Anywhere 12.0.1 build 3336 or newer. Query the `SYS.SYSHISTORY` system table to determine which version and build of the database server were used to initialize the database. If the database was initialized with a version older than 12.0.1 build 3336, update the database by executing `ALTER DATABASE UPGRADE PROCEDURE ON`.

You need to decide whether you want to run a separate database server, or whether you want to use the existing consolidated database as your message server. Consider the following when making this decision:

1. When a remote database authenticates with the message server, it uses the publisher of the remote database and the provided password to authenticate. While the user exists in the consolidated database, it

may not have a defined password (the remote user may not have connect privileges), which is a requirement of the HTTP message system. If granting CONNECT privilege to the remote users in the consolidated database is a security concern, set up a separate database to act as your message server.

2. If the consolidated database is heavily loaded, adding message server functionality to it may overwhelm the resources on the computer when the consolidated database runs.

To set up the required database objects for the database to act as a message server, call the `sr_add_message_server` stored procedure, which queries the SQL Remote definitions in the database.

If you are creating the message server as a separate database, you need to define a second database with SQL Remote definitions matching those of the consolidated database. Use the `dbunload` utility to create a copy of the consolidated database and specify the `-n` option to unload only the schema of the consolidated database, not the data:

```
dbunload -n -an -c "ServerName=cons,DBN=cons;UID=DBA;PWD=passwd"
```

If you are using a separate database as the message server, when changes are made to the SQL Remote definitions in the consolidated database, corresponding changes must also be made in the message server database.

To set up the message server, the directory where SQL Remote messages are stored must be accessible to the database server. To define the directory where messages are stored, use the `SET REMOTE OPTION` command and set the `root_directory` HTTP message parameter to the directory under which SQL Remote messages are stored. Next, choose the database user that will own the new objects that will be created, and ensure that the user is a role. Finally, execute the `sr_add_message_server` stored procedure, and pass in the name of the user that will own the objects.

Whenever changes are made to the SQL Remote definition of the message server (such as adding or removing remote users), run the `sr_update_message_system` stored procedure to update the definition of the objects required to support the message server. The message server will be unavailable for replication for a short period of time while the stored procedure runs and objects are dropped and recreated.

If you are no longer using the database as a message server, you can run the `sr_drop_message_system` stored procedure to remove the objects that were created to support the message server.

## Starting the Message Server Database Server

After the objects required to support the message server are created, when you start the message server database server you need to enable HTTP (and/or HTTPS) support to the database server using the `-xs` option.

The HTTP server-side protocol options of interest to those who have defined the objects needed for the message server:

### ServerPort | PORT

Specifies the port number that the database server uses to listen for HTTP or HTTPS requests in case the default ports of 80 and 443 are already being used on the computer.

### MaxRequestSize | MAXSIZE

Specifies the maximum size of a single HTTP request. The default value is 100 KB. If you have defined your SQL Remote messages size (`-l` option on the `dbremote` command line) to be greater than 100 KB, you also need to increase the size of the largest HTTP request that the database server can accept. The default SQL Remote message size is 50 KB.

### Identity (HTTPS only)

When you use HTTPS, the identity file contains the public certificate and its private key, and for certificates that are not self-signed, the identity file also contains all of the signing certificates, which include, among other things, the encryption certificate. The password for this certificate must be specified with the *Identity\_Password* parameter.

### Identity\_Password (HTTPS only)

When you use transport layer security, this option specifies the password that matches the password for the encryption certificate specified by the Identity protocol option.

## HTTP addresses and user IDs

To use SQL Remote and HTTP, each database participating in the system requires an HTTP address, a user ID and a password. These are distinct identifiers: the HTTP address is the destination of each message, and the user ID and password are the name and password entered by a user when they authenticate against the server.

## HTTP message control parameters

Before the SQL Remote Message Agent (dbremote) connects to the message system to send or receive messages, the user must have a set of control parameters already set on their computer, or the user is prompted to specify needed information. This information is needed only on the first connection. It is saved and used as the default for subsequent connections.

The HTTP message system uses the following control parameters that are set using the SET REMOTE OPTION statement:

### certificate

To make a secure (HTTPS) request, a client must have access to the certificate used by the HTTPS server. The necessary information is specified in a string of semicolon-separated keyword=value pairs. You can use the file keyword to specify the file name of the certificate. You cannot specify a file and certificate keyword together. The following keywords are available:

Keyword	Abbreviation	Description
file		The file name of the certificate
certificate	cert	The certificate itself
company	co	The company specified in the certificate
unit		The company unit specified in the certificate
name		The common name specified in the certificate

Certificates are required only for requests that are either directed to an HTTPS server, or can be redirected from a non-secure to a secure server. Only PEM formatted certificates are supported. `certificate='file=filename'`

### **client\_port**

Identifies the port number on which SQL Remote communicates using HTTP. It is provided for, and recommended only for, connections through firewalls that filter "outgoing" TCP/IP connections. You can specify a single port number, ranges of port numbers, or a combination of the two. Specifying a low number of client ports could result in SQL Remote being unable to send and receive messages if the operating system has not released the ports in a timely manner after SQL Remote closes the port on a previous run. `client_port= nnnnn [-mmmmm]`

### **debug**

When set to YES, all HTTP commands and responses are displayed in the output log. This information can be used for troubleshooting HTTP support problems. The default is NO.

### **https**

Specify whether to use HTTPS (`https=yes`) or HTTP (`https=no`).

### **password**

The message server database password. Authenticates to third-party HTTP servers and gateways using RFC 2617 Basic authentication. `password='password'`

### **proxy\_host**

Specifies the URI of a proxy server. For use when SQL Remote must access the network through a proxy server. Indicates that SQL Remote is to connect to the proxy server and send the request to the message server through it. `proxy_host='http://proxy-server[:port-number]'`

### **reconnect\_retries**

The number of times the link should try to open a socket with the server before failing. The default value is 4. When you set this parameter, only reconnections are affected. The initial connection made by the FTP link is not affected.

### **reconnect\_pause**

The time in seconds to pause between each connection attempt. The default setting is 30 seconds. When you set this parameter, only reconnections are affected. The initial connection made by the FTP link is not affected.

### **root\_directory**

This HTTP control parameter is ignored when specified at the client side. You define this control parameter in the message server prior to calling the `sr_add_message_server` or `sr_update_message_server` stored procedure. Specify the directory accessible by the message server under which the SQL Remote messages are stored. When using the HTTP message system, the address specified for a remote user or publisher can only contain a single subdirectory, and not multiple subdirectories. `root_directory='c:\msgs'`

### **url**

Specify the server name or IP address and, optionally, the port number of the HTTP server being used, separated by a semicolon. If requests are being passed through the Relay Server, you can optionally add a URL extension to indicate which server farm the request should be passed to. `url='server-name[:port-number][url-extension]'`

### **user**

The message server database user ID. Authenticates to third-party HTTP servers and gateways using RFC 2617 Basic authentication. `user='userid'`

## Related Information

[The FILE Message System \[page 125\]](#)

[The SMTP Message System \[page 134\]](#)

[The FTP Message System \[page 127\]](#)

[Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server \[page 190\]](#)

[Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server \[page 179\]](#)

[Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server \[page 201\]](#)

[CREATE REMOTE \[MESSAGE\] Statement \[SQL Remote\]](#)

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

[sr\\_add\\_message\\_server System Procedure \[page 239\]](#)

[sr\\_update\\_message\\_server System Procedure \[page 240\]](#)

[sr\\_drop\\_message\\_server System Procedure \[page 240\]](#)

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

[sr\\_add\\_message\\_server System Procedure \[page 239\]](#)

[-xs Database Server Option](#)

[ServerPort \(PORT\) Protocol Option](#)

[MaxRequestSize \(MAXSIZE\) Protocol Option](#)

[identity Protocol Option](#)

[identity\\_password Protocol Option](#)

### 1.4.7.9 The SMTP Message System

With the SMTP system, SQL Remote sends messages using Internet mail.

Messages are encoded in a text format and sent in an email message to the target database. The messages are sent using an SMTP server, and retrieved from a POP server.

#### SMTP Addresses and User IDs

To use SQL Remote and an SMTP message system, each database participating in the system requires an SMTP address and a POP3 user ID and password. These are distinct identifiers: the SMTP address is the destination of each message, and the POP3 user ID and password are the name and password entered by a user when they connect to their email server.

## i Note

Use a separate POP email account to send and receive SQL Remote messages.

## SMTP Message Control Parameters

Before the SQL Remote Message Agent (dbremote) connects to the message system to send or receive messages, the user must have a set of control parameters already set on their computer, or the user is prompted to specify needed information. This information is needed only on the first connection. It is saved in the ISYSREMOTEOPTION system table (see SYSREMOTEOPTION system view) and used as the default entries for subsequent connections.

The SMTP message system uses the following control parameters that are set by the SET REMOTE OPTION statement:

### **local\_host**

The name of the local computer. This is useful on computers where SQL Remote is unable to determine the local host name. The local host name is needed to initiate a session with any SMTP server. In most network environments, the local host name can be determined automatically and this entry is not needed.

### **pop3\_host**

The name of the computer on which the POP host is running. Typically, it is the same name as the SMTP host. This corresponds to the POP3 host field in the SMTP/POP3 login window.

### **pop3\_port**

The number of the port on which the POP server is listening. The default is 110.

### **pop3\_userid**

The user ID used to retrieve mail. The POP user ID corresponds to the user ID field in the SMTP/POP3 login window. You must obtain a user ID from your POP host administrator.

### **pop3\_password**

The password used to retrieve mail. This corresponds to the password field in the SMTP/POP3 login window.

### **smtp\_authenticate**

Determines whether the SMTP link authenticates the user. The default value is YES. Set this parameter to NO to turn off SMTP authentication.

### **smtp\_host**

The name of the computer on which the SMTP server is running. This corresponds to the SMTP host field in the SMTP/POP3 login window.

### **smtp\_port**

The number of the port on which the SMTP server is listening. The default is 25.

### **smtp\_userid**

The user ID for SMTP authentication. By default, this parameter takes the same value as the pop3\_userid parameter. The smtp\_userid only needs to be set if the user ID is different from that of the POP server.

### **smtp\_password**

The password for SMTP authentication. By default, this parameter takes the same value as the `pop3_password` parameter. The `smtp_password` only needs to be set if the user ID is different from that of the POP server.

#### **suppress\_dialogs**

This parameter is set to ON or OFF. The default value is OFF. When this parameter is set to ON, the [Connect](#) window does not appear after failed attempts to connect to the mail server. Instead, an error is generated.

#### **top\_supported**

SQL Remote uses a POP3 command called TOP when enumerating incoming messages. The TOP command may not be supported by all POP servers. When you set the `top_supported` parameter to NO, SQL Remote uses the RETR command, which is less efficient but works with all POP servers. The default is YES.

#### **debug**

When set to YES, all SMTP and POP3 commands and responses are displayed in the output log. This information can be used for troubleshooting SMTP/POP support problems. The default is NO.

#### **encode\_dll**

If you have implemented a custom encoding scheme, you must set this to the full path of the custom encoding DLL that you created.

#### **max\_retries**

By default, when SQL Remote is running in continuous mode and an error occurs when accessing the message system, it shuts down after the send and/or received phases. Use this parameter to specify the number of times you want SQL Remote to retry the send and/or receive phases before it shuts down.

#### **output\_log\_send\_limit**

Limits the amount of information that is sent to the consolidated database. The `output_log_send_limit` option specifies the number of bytes at the end of the output log (that is, the most recent entries) that are sent to the consolidated database. The default is 5K.

#### **output\_log\_send\_on\_error**

Sends log information when an error occurs.

#### **output\_log\_send\_now**

When set to YES, sends output log information to the consolidated database. On the next poll, the remote database sends the output log information and then resets the `output_log_send_now` option to NO.

#### **pause\_after\_failure**

This parameter applies when the `max_retries` parameter is specified to a value other than zero and SQL Remote is running in continuous mode. When an error occurs in the message system, this parameter defines the number of seconds SQL Remote waits between retrying the send and/or receive phases.

### **In this section:**

#### [SMTP/POP Address Sharing \[page 137\]](#)

You should use a separate email account for SQL Remote messages. It is not recommended that you send and receive SQL Remote messages through the same email account that you use for personal or business email messages.

#### [Troubleshooting the SMTP Link \[page 137\]](#)



If you cannot get the SMTP Link to work, connect to the SMTP/POP3 server from the same computer on which the SQL Remote Message Agent (dbremote) is running, using the same account and password.

## Related Information

[The FILE Message System \[page 125\]](#)

[The FTP Message System \[page 127\]](#)

[Message Size \[page 118\]](#)

[Supported Platforms](#)

[CREATE REMOTE \[MESSAGE\] Statement \[SQL Remote\]](#)

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

### 1.4.7.9.1 SMTP/POP Address Sharing

You should use a separate email account for SQL Remote messages. It is not recommended that you send and receive SQL Remote messages through the same email account that you use for personal or business email messages.

If you need to share the same email account for SQL Remote messages and regular email messages, then you must ensure that your email program does not download and delete all messages from the mail server, including SQL Remote email and personal messages. You must configure the email program so that it does not alter or delete SQL Remote messages when it downloads your regular email messages. SQL Remote messages contain the subject ---SQL Remote--.

### 1.4.7.9.2 Troubleshooting the SMTP Link

If you cannot get the SMTP Link to work, connect to the SMTP/POP3 server from the same computer on which the SQL Remote Message Agent (dbremote) is running, using the same account and password.

Use an Internet email program that supports SMTP/POP3. Disable this program once the SMTP message link is working.

## Check that Email is Working Properly

If SQL Remote messages are not being sent and received properly and you are using an email message system, you should confirm that email is working properly between the two computers.

## 1.4.8 SQL Remote System Backups

There are a number of considerations to be aware of before you perform backup activities.

SQL Remote replication depends on access to operations in the transaction log, and access to old transaction logs. Any backup strategy that you implement must incorporate the maintenance of the transaction logs for SQL Remote.

To perform backup activities you require the BACKUP DATABASE system privilege.

The SQL Remote Message Agent (dbremote) must have access to the current and old transaction logs until they are no longer needed.

A consolidated database no longer needs its transaction logs when all remote databases have received and have confirmed that the messages contained in the transaction logs have been successfully applied.

A remote database no longer needs its transaction logs when the consolidated database has received and confirmed that it has successfully applied the messages contained in the transaction logs.

### Backing up Remote Databases

For your remote databases, you need to decide whether to:

#### **Rely on replication to the consolidated database as a backup method**

Backup procedures are not as crucial on remote databases as on the consolidated database. You can rely on replication to the consolidated database as a data backup method.

If you choose this method, then you *should* create a strategy for maintaining the remote database transaction logs.

#### **Create a backup strategy for the remote database**

If the changes made on the remote databases are crucial, then you need to create a backup strategy for the remote databases that includes the maintenance of the transaction logs.

### Backing up Consolidated Databases

You *must* have a backup strategy for your consolidated database that includes the maintenance of the transaction logs.

### The Backup Utility (dbbackup) and the SQL Remote Message Agent (dbremote) -x Option

On a database, you should never run both the SQL Remote Message Agent (dbremote) with the -x option and the Backup utility (dbbackup).

The -x option is used to manage transaction logs for replication. The -x option ensures that SQL Remote Message Agent has access to old transaction logs and deletes the transaction logs when they are no longer needed. The -x option does not back up the transaction log.

The Backup utility (dbbackup) is used to back up the current transaction log. When the Backup utility (dbbackup) is run with the -r and -n options, it backs up the current transaction log to a backup directory and renames and restarts the current transaction log. The Backup utility (dbbackup) assumes that the current transaction log is the same transaction log that it renamed and restarted after the previous back up.

If you try to run both the SQL Remote Message Agent -x option and the Backup utility (dbbackup) on the same database, they interfere with each other. You can lose transaction logs when both are running.

Only run the SQL Remote Message Agent (dbremote) with the -x option on a remote database that is not being backed up.

#### **In this section:**

##### [Maintaining Transaction Logs for Remote Databases \(Command Line\) \[page 139\]](#)

Maintain remote database transaction logs when you are relying on replication to the consolidated database to back up your remote databases. That is, you are *not* running the Backup utility (dbbackup) on the remote databases and transaction logs.

##### [Protect the Consolidated Database from Media Failures \[page 142\]](#)

SQL Remote supports a number of strategies for reducing the risk of data loss due to media failure.

##### [Backing Up the Consolidated Database \(Command Line\) \[page 143\]](#)

Back up your consolidated database by making a full backup of the consolidated database and transaction log, and then make incremental backups of the transaction log.

## **Related Information**

[Backing Up a Remote Database \(Command Line\) \[page 141\]](#)

### **1.4.8.1 Maintaining Transaction Logs for Remote Databases (Command Line)**

Maintain remote database transaction logs when you are relying on replication to the consolidated database to back up your remote databases. That is, you are *not* running the Backup utility (dbbackup) on the remote databases and transaction logs.

## **Context**

#### **⚠ Caution**

Do *not* run the SQL Remote Message Agent (dbremote) with the -x option on a database that is being backed up.

## Procedure

1. On the remote database, run the SQL Remote Message Agent (dbremote) with the -x option and specify a size for the transaction log. This option causes the SQL Remote Message Agent (dbremote) to rename and restart the transaction log when the transaction log exceeds the specified size.

The following command deletes the transaction log when it is larger than 1 MB:

```
dbremote -x 1M -c "UID=ManagerSteve;PWD=passwd;DBF=c:\mydata.db"
```

2. On the remote database, set the delete\_old\_logs option to On. Setting the delete\_old\_logs\_option causes the old transaction log files to be deleted automatically by dbremote when they are no longer needed for replication.

A transaction log is no longer needed when all subscribers have confirmed that they have received and successfully applied all the changes recorded in that transaction log file. You can set the delete\_old\_logs option either for the PUBLIC role or just for the user contained in the dbremote connection string.

The following statement sets the public delete\_old\_logs option to delete transaction logs that were created more than 10 days ago:

```
SET OPTION PUBLIC.delete_old_logs = '10 days';
```

## Results

The database transaction logs are deleted according to the specified rules.

### In this section:

[Backing Up a Remote Database \(Command Line\) \[page 141\]](#)

Back up your remote databases.

## Related Information

[delete\\_old\\_logs Option \[MobiLink\]\[SQL Remote\]](#)

## 1.4.8.1.1 Backing Up a Remote Database (Command Line)

Back up your remote databases.

### Prerequisites

You must have the DATABASE BACKUP system privilege. This procedure includes a maintenance strategy for the use of the transaction logs by SQL Remote. Do *not* use this procedure and run the SQL Remote Message Agent (dbremote) with the -x option.

### Procedure

1. Make a full backup of the remote database.
  - a. Connect to the database.
  - b. Run dbbackup with the -r and -n options.

For example, assume that the backup directory is `e:\backup_dir`, the database file is located in the `c:\live` directory, and its corresponding transaction log file is located in the `d:\live` directory:

```
dbbackup -r -n -c "UID=DBA;PWD=passwd;DBF=c:\live\remotedatabase.db" e:\backup_dir
```

The transaction logs in the `d:\live` directory are not altered by the full backup.

- c. Copy the backup files located in the `e:\backup_dir` directory to an off-site drive or to a DVD.
- d. Run the SQL Remote Message Agent (dbremote) with access to the current transaction log files using the following command:

```
dbremote -c "UID=DBA;PWD=passwd;DBF=c:\live\remotedatabase.db" d:\live
```

#### Caution

Do *not* run the SQL Remote Message Agent (dbremote) with the -x option on a database that is being backed up.

2. Set up the Backup utility (dbbackup) to make incremental backups of the remote database's transaction log.
  - a. Connect to the database.
  - b. Run dbbackup with the -r, -n, and -t options.

For example:

```
dbbackup -r -n -t -c "UID=DBA;PWD=passwd;DBF=c:\live\remotedatabase.db" e:\backup_dir
```

- c. Run the SQL Remote Message Agent (dbremote) with access to the current transaction log files using the following command:

```
dbremote -c "UID=DBA;PWD=passwd;DBF=c:\live\remotedatabase.db" d:\live
```

## Results

The remote database is backed up.

### 1.4.8.2 Protect the Consolidated Database from Media Failures

SQL Remote supports a number of strategies for reducing the risk of data loss due to media failure.

To protect your SQL Remote replication system against media failure:

#### Replicate only backed-up transactions

Send messages that contain only backed-up transactions. By sending only backed-up transactions, the replication system is protected against media failure on the transaction log. You can accomplish this by:

- Running the SQL Remote Message Agent (dbremote) with the -u option. When the SQL Remote Message Agent (dbremote) is run with the -u option, only committed transactions that have been backed up are packaged into messages to be sent.  
The -u option provides additional protection against total site failure, if backups are carried out to another site.

#### Use a transaction log mirror

Using a transaction log mirror protects against media failure on the transaction log device.

#### Do not run SQL Remote Message Agent (dbremote) with the -x option on the consolidated database

Never run the SQL Remote Message Agent (dbremote) with the -x option against a database that is being backed up. The -x option maintains the transaction logs for replication, not for backup or recovery.

## Related Information

[Transaction Log Mirrors](#)

[Backing Up the Consolidated Database \(Command Line\) \[page 143\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

## 1.4.8.3 Backing Up the Consolidated Database (Command Line)

Back up your consolidated database by making a full backup of the consolidated database and transaction log, and then make incremental backups of the transaction log.

### Prerequisites

You must have the BACKUP DATABASE system privilege.

### Context

#### ⚠ Caution

Do *not* run the SQL Remote Message Agent (dbremote) with the -x option on a database that is being backed up.

### Procedure

1. Make a full back up of the consolidated database and its transaction log.
  - a. Connect to the database.
  - b. Run dbbackup with the -r and -n options.

For example:

```
dbbackup -r -n -c "UID=DBA;PWD=passwd;DBF=c:\live\database.db" e:\archive
```

2. Make incremental backups of the consolidated database transaction log. When backing up the transaction log, choose to rename and restart the transaction log.
  - a. Connect to the database.
  - b. Run dbbackup with the -r, -n and -t options.

For example:

```
dbbackup -r -n -t -c "UID=DBA;PWD=passwd;DBF=c:\live\database.db" e:\archive
```

3. Run the SQL Remote Message Agent (dbremote) with access to the current transaction log.

For example:

```
dbremote -c "UID=DBA;PWD=passwd;DBF=c:\live\database.db" d:\live
```

## Results

The consolidated database and transaction log are backed up.

## Example

Consider a database named `database.db` in the `c:\live` directory, with a transaction log named `database.log` in the `d:\live` directory.


When you back up the transaction log to the backup directory `e:\archive` using the `-r` and `-n` options to rename and restart the transaction log, the Backup utility (`dbbackup`) carries out the following tasks:

1. Renames the current transaction log file to `971201xx.log`, where `xx` are sequential characters ranging from `AA` to `ZZ`.
2. Backs up the transaction log file to the backup directory, creating a backup file named `971201xx.log`.

### Note

Before the release of SQL Anywhere 8.0.1, the old transaction log files were named `yyymmdd01.log`, `yyymmdd02.log`, and so on. The name change was introduced to allow more old transaction logs to be stored. As the SQL Remote Message Agent (`dbremote`) scans all the files in the specified directory, regardless of their names, the name change should not affect existing applications.

3. Starts a new transaction log, as `database.log`.  
After several backups, the live directory and the archive directory contain a set of sequential transaction logs.



## Related Information

[Renaming or Truncating the Transaction Log \(SQL Central\)](#)

[Protect the Consolidated Database from Media Failures \[page 142\]](#)

[Backup Utility \(dbbackup\)](#)



## 1.4.9 Recovering a Consolidated Database Manually (Command Line)

Recover a consolidated database by applying each transaction log to the database.

### Prerequisites

You must have the BACKUP DATABASE system privilege.

Make a copy of the database and transaction log file. This procedure assumes that previous backups of the database file have been made and are available.

### Procedure

1. Create a temporary directory.
2. Restore the most recent back up of the database (.db) file, *not* the transaction log file.

In the temporary directory:

- a. Start the backup copy of the database.
  - b. Apply the old transaction logs using the -a option.
  - c. Shut down the database.
  - d. Start the database using the current transaction log and the -a option to apply the transactions and bring the database file up to date.
  - e. Shut down the database.
  - f. Back up the database.
3. Copy the database to the production directory.
  4. Start the database.

Any new activity is appended to the current transaction log.

### Results

The consolidated database is recovered.

### Example

Suppose you have a consolidated database file named `c:\dbdir\cons.db`, a transaction log file `c:\dbdir\cons.log`, and a transaction log mirror file `d:\mirmdir\cons.mlg`.

Assume that you perform full backups weekly, and you perform incremental backups daily using the following command:

```
dbbackup -c "UID=DBA;PWD=passwd" -r -n -t e:\backdir
```

This command backs up the transaction log `cons.log` to the directory `e:\backdir`. The transaction log file is then renamed to `datexx.log`, where `date` is the current date and `xx` is the next set of letters in sequence, and a new transaction log is started. The directory `e:\backdir` is then backed up using a third-party utility.

In this scenario, you run the SQL Remote Message Agent (`dbremote`) with the optional directory to point to the renamed transaction log files. For example:

```
dbremote -c "UID=DBA;PWD=passwd" c:\dbdir
```

On the third day following the weekly backup, the database file is corrupted because of a bad disk block. Perform the following steps:

1. Back up the transaction log mirror file `d:\mirmdir\cons.mlg`.
2. Create a temporary directory to perform the recovery in. In this example, the directory is called `c:\recover`.
3. Restore the most recent backup of the database file, `cons.db` to `c:\recover\cons.db`.

```
dbeng17 -a c:\dbdir\dateAA.log c:\recover\cons.db  
dbeng17 -a c:\dbdir\dateAB.log c:\recover\cons.db
```

4. Apply the renamed transaction logs in order, as follows:
5. Copy the current transaction log, `d:\mirmdir\cons.log` to the recovery directory, giving `c:\recover\cons.log`.

```
dbeng17 c:\recover\cons.db
```

6. Start the database using the following command:
7. Shut down the database server.
8. Back up the recovered database and transaction log from `c:\recover`.
9. Copy the files from `c:\recover` to the appropriate production directories:
  - Copy `c:\recover\cons.db` to `c:\dbdir\cons.db`.
  - Copy `c:\recover\cons.log` to `c:\dbdir\cons.log`, and to `d:\mirmdir\cons.mlg`.

## Related Information

[Recovering a Consolidated Database Automatically \(Command Line\) \[page 147\]](#)  
[-a Database Option](#)

## 1.4.10 Recovering a Consolidated Database Automatically (Command Line)

Recover a consolidated database automatically.

### Prerequisites

You must have the BACKUP DATABASE system privilege.

Make a copy of the database and transaction log file. This procedure assumes that previous backups of the database file have been made and are available.

### Procedure

1. Restore the most recent backed up copy of the database (.db) file, *not* the transaction log file, into a temporary directory.

2. In the temporary directory:

- a. Start the database, applying the transaction logs using the -ad option.

When you specify the -ad option, the database server looks in the specified directory for the transaction logs for the database. It then determines the correct order to apply the transaction logs based on the transaction log offsets.

- b. Copy the current transaction log to the temporary directory.
  - c. Start the database and apply the current transaction log.
  - d. Shut down the database server.
  - e. Back up the database and transaction log.
3. Copy the database and transaction log files to the appropriate production directories.
  4. Restart your system as normal.

Any new activity is appended to the current transaction log.

### Results

The consolidated database is recovered.

### Example

Suppose you have a consolidated database file named `c:\dbdir\cons.db`, a transaction log file `c:\dbdir\cons.log`, and a transaction log mirror file `d:\mirmdir\cons.mlg`.

Assume that you perform full backups weekly using the following command:

```
dbbackup -c "UID=DBA;PWD=passwd" -r -n e:\backdir
```

Assume that you also perform incremental backups daily using the following command:

```
dbbackup -c "UID=DBA;PWD=passwd" -r -n -t e:\backdir
```

This command backs up the transaction log `cons.log` to the directory `e:\backdir`. The transaction log file is then renamed to `datexx.log`, where `date` is the current date and `xx` is the next set of letters in sequence, and a new transaction log is started. The directory `e:\backdir` is then backed up using a third-party utility.

In this scenario, you would run the SQL Remote Message Agent (`dbremote`) with the optional directory to point to the renamed transaction log files. For example:

```
dbremote -c "UID=DBA;PWD=passwd" c:\dbdir
```

On the third day following the weekly backup, the database file is corrupted because of a bad disk block. Perform the following steps:

1. Replace the `c:\` drive.
2. Back up the transaction log mirror file `d:\mirdir\cons.mlg`.
3. Create a temporary directory to perform the recovery in. In this example, it is called `c:\recover`.
4. Restore the most recent backup of the database file, `cons.db` to `c:\recover\cons.db`.
5. Copy the backed up transaction logs to `c:\dbdir`.
6. Apply the renamed transaction logs:

```
dbeng17 c:\recover\cons.db -ad c:\dbdir
```

7. Copy the current transaction log, `d:\mirdir\cons.log` to the recovery directory, giving `c:\recover\cons.log`.
8. Start the database using the following command:

```
dbeng17 c:\recover\cons.db
```

9. Shut down the database server.
10. Back up the recovered database and transaction log from `c:\recover`.
11. Copy the files from `c:\recover` to the appropriate production directories:
  - Copy `c:\recover\cons.db` to `c:\dbdir\cons.db`.
  - Copy `c:\recover\cons.log` to `c:\dbdir\cons.log`, and to `d:\mirdir\cons.mlg`.

## Related Information

[Recovering a Consolidated Database Manually \(Command Line\) \[page 145\]](#)

[-ad Database Option](#)

[Backup Utility \(dbbackup\)](#)

## 1.4.11 Replication Error Reporting and Handling

SQL Remote provides a number of methods for handling potential errors.

SQL Remote systems may encounter the following errors:

- Row not found errors
- Referential integrity errors
- Duplicate primary key errors

By default, when an error occurs, the SQL Remote Message Agent (dbremote) prints the error in its log output window. The SQL Remote Message Agent (dbremote) can print more information in the output messages file than in the messages window.

The SQL Remote Message Agent (dbremote) messages log file includes the following information:

- Applied messages
- Failed SQL statements
- Other errors

To print an error to the output log file, run the SQL Remote Message Agent (dbremote) with the `-o` option.

When an error occurs, you can configure SQL Remote to:

### Run an error-handling procedure

By default, no procedure is called. However, you can use the `replication_error` database option to specify a stored procedure to be called by the SQL Remote Message Agent (dbremote) when an error occurs.

For example, you can configure SQL Remote to:

- Send portions of a remote database's output log to the consolidated database and written to a file.
- Send an email notification when an error occurs at a remote database.

### Ignore the error

There might be instances when you do not want the SQL Remote Message Agent (dbremote) to report an error. For example, you can choose to ignore an error when you know the conditions under which the error occurs and you are sure that the error does not produce inconsistent data.

### In this section:

[Replication Error-handling Procedures \[page 150\]](#)

Set the `replication_error` option to call a procedure when a SQL error occurs.

[How to Ignore Replication Errors \[page 154\]](#)

Create a BEFORE trigger on the action that causes the known error. This trigger should signal an error.

## Related Information

[Collecting Errors from the Remote Database \(SQL\) \[page 151\]](#)

[Receiving Email Notification about Remote Database Errors \(SQL\) \[page 152\]](#)

[Row Not Found Errors \[page 55\]](#)

[Referential Integrity Errors \[page 56\]](#)

[Duplicate Primary Key Errors \[page 58\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

## 1.4.11.1 Replication Error-handling Procedures

Set the `replication_error` option to call a procedure when a SQL error occurs.

By default, no procedure is called when a SQL error occurs.

The procedure that is called must have a single argument type of CHAR, VARCHAR, or LONG VARCHAR. The procedure is called once with the SQL error message and once with the SQL statement that causes the error.

To set the `replication_error` option, execute the following statement. `remote-user` is the publisher name on the SQL Remote Message Agent (dbremote) command and `procedure-name` is the procedure called when a SQL error is detected.

```
SET OPTION
remote-user.replication_error
= 'procedure-name';
```

### In this section:

#### [Collecting Errors from the Remote Database \(SQL\) \[page 151\]](#)

Send portions of a remote database's output log to the consolidated database. The information is written to a file and the file can contain output logging information from some or all remote databases in the system.

#### [Receiving Email Notification about Remote Database Errors \(SQL\) \[page 152\]](#)

Send email notifications when errors occur at a remote database. You can use email or a paging system to receive the notifications.

## Related Information

[replication\\_error Option \[SQL Remote\]](#)

## 1.4.11.1.1 Collecting Errors from the Remote Database (SQL)

Send portions of a remote database's output log to the consolidated database. The information is written to a file and the file can contain output logging information from some or all remote databases in the system.

### Procedure

1. Configure the remote databases to send output log information to the consolidated database.
  - a. Use the SET REMOTE statement with the output\_log\_send\_on\_error option to send log information when an error occurs.

On the remote database, execute the following statement:

```
SET REMOTE link-name OPTION  
PUBLIC.output_log_send_on_error = 'Yes';
```

When the SQL Remote Message Agent (dbremote) reads any messages that start with the error indicator E, it sends the output log information to the consolidated database.

- b. This step is optional. Set the SET REMOTE statement with the output\_log\_send\_limit option to limit the amount of information that is sent to the consolidated database. The output\_log\_send\_limit option specifies the number of bytes at the end of the output log (that is, the most recent entries) that are sent to the consolidated database. The default is 5K.

If you supply an output\_log\_send\_limit value that exceeds the maximum message size, SQL Remote overrides the output\_log\_send\_limit value and sends only what can fit within the maximum message size.

On the remote database, execute the following statement:

```
SET REMOTE link-name OPTION  
PUBLIC.output_log_send_limit = '7K';
```

2. Configure the consolidated database to receive log information.

On the consolidated database, run the SQL Remote Message Agent (dbremote) with either the -ro or the -rt options.

3. This step is optional. To test your configurations, set the output\_log\_send\_now option to send output log information to the consolidated database.

On the remote database, set the output\_log\_send\_now option to YES.

On the next poll, the remote database sends the output log information and then resets the output\_log\_send\_now option to NO.

### Results

Output log information is now sent to the consolidated database.

## Related Information

[Receiving Email Notification about Remote Database Errors \(SQL\) \[page 152\]](#)

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

### 1.4.11.1.2 Receiving Email Notification about Remote Database Errors (SQL)

Send email notifications when errors occur at a remote database. You can use email or a paging system to receive the notifications.

#### Procedure

1. Connect to the consolidated database as user Cons.
2. Create a stored procedure that notifies the DBA user by email that an error has occurred.

For example, execute the following to create the sp\_LogReplicationError procedure:

```
CREATE PROCEDURE cons.sp_LogReplicationError
  ( IN error_text LONG VARCHAR )
BEGIN
  DECLARE current_remote_user CHAR( 255 );
  SET current_remote_user = CURRENT REMOTE USER;
  // Log the error
  INSERT INTO cons.replication_audit
    ( remoteuser, errormsg )
  VALUES
    ( current_remote_user, error_text );
  COMMIT WORK;
  //Now notify the DBA by email that an error has occurred
  // on the consolidated database. The email should contain the error
  // strings that the SQL Remote Message Agent is passing to the procedure.
  IF CURRENT PUBLISHER = 'cons' THEN
    CALL sp_notify_DBA( error_text );
  END IF
END;
```

3. Create a stored procedure that manages the sending of email.

For example, execute the following statement to create the sp\_notify\_DBA procedure:

```
CREATE PROCEDURE sp_notify_DBA( in msg long varchar)
BEGIN
  DECLARE rc INTEGER;
  rc=call xp_startmail( mail_user='davidf' );
  //If successful logon to mail
  IF rc=0 THEN
  rc=call xp_sendmail(
    recipient='Doe, John; Smith, Elton',
    subject='SQL Remote Error',
    "message"=msg);
  //If mail sent successfully, stop
```



```

IF rc=0 THEN
    call xp_stopmail()
END IF
END IF
END;

```

4. Set the replication\_error database option to call the procedure that notifies the DBA by email that an error occurs.

For example, execute the following statement to call the sp\_LogReplicationError procedure when an error occurs:

```

SET OPTION PUBLIC.replication_error =
'cons.sp_LogReplicationError';

```

5. Create an audit table.

For example, execute the following to create the replication\_audit table:

```

CREATE TABLE replication_audit (
    id    INTEGER DEFAULT AUTOINCREMENT,
    pub   CHAR(30)  DEFAULT CURRENT PUBLISHER,
    remoteuser CHAR(30),
    errormsg LONG VARCHAR,
    timestamp DATETIME DEFAULT CURRENT TIMESTAMP,
    PRIMARY KEY (id,pub)
);

```

The following table describes the columns of the replication\_audit table:

Column	Description
pub	Current publisher of the database (identifies the database in which the publisher was inserted).
remoteuser	Remote user applying the message (identifies the database from which the remote user came from).
errormsg	Error message passed to the replication_error procedure.

6. Test your procedures.

For example, insert a row on the consolidated database that uses the same primary key as a row on a remote database. When this row from the consolidated database is replicated to the remote database, a primary key conflict error occurs and:

- The remote database SQL Remote Message Agent (dbremote) prints the following message to its output log:

```

Received message from "cons" (0-0000000000-0)
SQL statement failed: (-193) primary key for table 'reptable' is not unique
INSERT INTO cons.reptable( id,text,last_contact )
VALUES (2,'dave','1997/apr/21 16:02:38.325')
COMMIT WORK

```

- The following INSERT statement is sent to the consolidated database:

```

INSERT INTO cons.replication_audit
( id,
  pub,
  remoteuser,
  errormsg,
  "timestamp")

```

```
VALUES
( 1,
  'cons',
  'sales',
  'primary key for table ''reptable'' is not unique (-193)',
  '1997/apr/21 16:03:13.836');
COMMIT WORK;
```

- An email is sent to John Doe and Elton Smith with the following message:

```
primary key for table 'reptable' is not unique (-193)
INSERT INTO cons.reptable( id,text,last_contact )
VALUES (2,'dave','1997/apr/21 16:02:52.605')
```

## Results

Email notification are now sent when errors occur at a remote database.

## Related Information

[Collecting Errors from the Remote Database \(SQL\) \[page 151\]](#)

### 1.4.11.2 How to Ignore Replication Errors

Create a BEFORE trigger on the action that causes the known error. This trigger should signal an error.

For example, to ignore INSERT statement errors that occur when a table is missing a referenced column, create a BEFORE INSERT trigger that signals the `SQL_REMOTE_STATEMENT_FAILED` SQLSTATE when the referenced column does not exist. The INSERT statement fails, but this failure is not reported in the SQL Remote Message Agent (dbremote) output log.

## Related Information

[Remote statement failed](#)

### 1.4.12 SQL Remote Security Considerations

SQL Remote provides a number of features to protect your data.

SQL Remote provides the following security features:

**SYS\_RUN\_REPLICATION\_ROLE** system role

Connect to the SQL Remote Message Agent (dbremote) with a user that has the SYS\_RUN\_REPLICATION\_ROLE system role.

### **Database encryption**

You can encrypt your database using the -ek option.

### **Message encryption**

The SQL Remote Message Agent (dbremote) uses a simple obfuscation algorithm to protect the messages against casual snooping. However, this encryption scheme is not intended to provide full protection against determined efforts to decipher them. You can implement your own custom encoding scheme if security of your data over the network is an issue.

## **Related Information**

[Database Encryption and Decryption](#)

[GRANT ROLE Statement](#)

[Extraction Utility \(dbxtract\) \[page 225\]](#)

## **1.4.13 Upgrades and Resynchronization**

There are a number of different ways to upgrade and resynchronize SQL Remote databases.

You must have the SYS\_RUN\_REPLICATION\_ROLE system role.

Use caution when upgrading a SQL Remote system. You can upgrade a SQL Remote system in any of the following ways:

### **Upgrading software**

Upgrade the software.

### **Changing the database schemas**

To make changes to the database schema, you can:

#### **Use passthrough mode**

The passthrough mode allows schema changes to be sent to some or all databases in a SQL Remote system, but it requires careful planning and execution.

#### **Re-synchronize subscriptions**

Re-synchronization involves copying new copies of the data to the remote databases. When there are many remote databases, resynchronization can be a time-consuming process involving work interruptions and possible data loss.

### **In this section:**

[Changes to Avoid on a Running System \[page 156\]](#)

There are a few limitations to be aware of if you are making changes to deployed and running SQL Remote systems.

## Related Information

[SQL Remote Passthrough Mode \[page 157\]](#)

[Subscription Resynchronization \[page 160\]](#)

[SQL Remote Upgrades](#)

[PASSTHROUGH Statement \[SQL Remote\]](#)

### 1.4.13.1 Changes to Avoid on a Running System

There are a few limitations to be aware of if you are making changes to deployed and running SQL Remote systems.

The following changes should not be made to a deployed and running SQL Remote system, except under the conditions stated:

#### Changing publishers

Problems can occur if you change the publisher user name on a consolidated database of a deployed and running SQL Remote system. If you need to change the consolidated database publisher user name, you must shut down the SQL Remote system and resynchronize all remote users.

Changing the user name of a publisher at a remote database causes problems for any subscriptions that the remote database is involved in, including the loss of information. If you need to change a remote database publisher user name, shut down the remote database and resynchronize the remote user.

#### Making restrictive changes to tables

You cannot make restrictive changes to tables. For example, do not drop a column or alter a column to disallow NULL values because messages can exist in the system that reference these columns.

#### Altering publications

Publication definitions must be maintained on both the consolidated and the remote databases. Altering publications in a running SQL Remote system can cause replication errors and can lead to a loss of data in the replication system.

#### Deleting subscriptions

You can delete a subscription, but you must use passthrough mode to remove the data on the remote database.

#### Unloading and reloading databases

You must ensure that the transaction log is properly maintained.

#### Making changes in a multi-tier hierarchy

You make changes to multi-tier hierarchy systems using the dbxtract or *Extract Database Wizard*.

## Related Information

[SQL Remote Passthrough Mode \[page 157\]](#)

[Passthrough Mode Limitations \[page 158\]](#)

[Subscription Resynchronization \[page 160\]](#)

[Rebuilding Databases Involved in Synchronization or Replication \(dbunload\)](#)

[ALTER PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)

[PASSTHROUGH Statement \[SQL Remote\]](#)

[Database Extraction for a Multi-tier Hierarchy System \[page 92\]](#)

## 1.4.14 SQL Remote Passthrough Mode

Use passthrough mode to pass standard SQL statements to a remote database where they can be executed.

You must have the SYS\_RUN\_REPLICATION\_ROLE system role.

You can use passthrough mode to complete the following tasks on a running SQL Remote system:

- Add new users.
- Resynchronize users.
- Delete users from the system.
- Change the address, message type, or frequency for a remote user.
- Add a column to a table.

### Caution

- SQL Remote relies on each database in the system having the same objects; when a table is altered at some sites but not at others, attempts to replicate data changes fail. Additional schema changes executed on a running SQL Remote system might cause problems.
- Test your passthrough operations on a copy of the consolidated database with a copy of a remote database subscribed. Never run untested passthrough scripts on a production database.
- Qualify object names with the owner name. PASSTHROUGH statements are not executed on remote databases from the same user name. Object names without the owner name qualifier may not be resolved correctly.

### In this section:

[Passthrough Mode Limitations \[page 158\]](#)

Passthrough mode has a number of limitations to be aware of.

[Start and Stop Passthrough Mode \[page 159\]](#)

Passthrough mode is started using the PASSTHROUGH statement and it is stopped using the PASSTHROUGH STOP statement. A passthrough session refers to the statements entered between the PASSTHROUGH statements.

## Related Information

[Changes to Avoid on a Running System \[page 156\]](#)

[PASSTHROUGH Statement \[SQL Remote\]](#)

## 1.4.14.1 Passthrough Mode Limitations

Passthrough mode has a number of limitations to be aware of.

### **Passthrough works on only one level of a hierarchy**

In a multi-tier SQL Remote system, it is important that passthrough statements work immediately below the current level. In a multi-tier system, passthrough statements must be entered at each consolidated database, for the level beneath it.

### **Calling procedures**

When a stored procedure is called in passthrough mode using a CALL or EXEC statement:

- The procedure must exist in the consolidated database that calls the passthrough command, even if the procedure is not executed on the consolidated database.
- The procedure must also exist on the remote database. The CALL or EXEC statement is replicated, but none of the statements inside the procedure is replicated. It is assumed that the procedure on the replicated database has the correct effect.

### **Control statements**

Control statements such as IF and LOOP and any cursor operations, are not replicated in passthrough mode. Any statements within the loop or control structure *are* replicated.

### **Cursor operations**

Operations on cursors are not replicated.

### **SQL SET OPTION statements**

Static Embedded SQL SET OPTION statements are not replicated. However, dynamic SQL statements are replicated.

For example, the following statement is not replicated in passthrough mode:

```
EXEC SQL SET OPTION ...
```

However, the following dynamic SQL statement is replicated:

```
EXEC SQL EXECUTE IMMEDIATE "SET OPTION ... "
```

### **Batch statements**

Batch statements (a group of statements surrounded with a BEGIN and END) are not replicated in passthrough mode. If you try to use batch statements in passthrough mode, an error occurs.

## Related Information

[Control Statements](#)

[Subscription Resynchronization \[page 160\]](#)

[Static and Dynamic SQL](#)

[PASSTHROUGH Statement \[SQL Remote\]](#)

## 1.4.14.2 Start and Stop Passthrough Mode

Passthrough mode is started using the PASSTHROUGH statement and it is stopped using the PASSTHROUGH STOP statement. A passthrough session refers to the statements entered between the PASSTHROUGH statements.

Statements entered in a passthrough session:

- Are checked for syntax errors.
- Are executed at the consolidated database unless you supply the ONLY keyword. When ONLY is specified, the statements are sent to the remote database without being executed on the consolidated database. The following statement starts a passthrough session, which passes the statements to a list of two named subscribers, without being executed at the current database:

```
PASSTHROUGH ONLY  
FOR userid_1, userid_2;
```

- Are passed to the identified subscriber database. Passthrough statements are replicated in sequence with normal replication messages, in the order in which the statements are recorded in the transaction log.
- Are executed at the subscriber database.

### Direct Passthrough Statements

The following statement starts a passthrough session that passes the statements to all users who are subscribed to the pubname publication:

```
PASSTHROUGH ONLY  
FOR SUBSCRIPTION TO [owner].pubname statement1;
```

Passthrough mode is additive. In the following example, statement\_1 is sent to user\_1, and statement\_2 is sent to both user\_1 and user\_2.

```
PASSTHROUGH ONLY FOR user_1 ;  
statement_1;  
PASSTHROUGH ONLY FOR user_2 ;  
statement_2;
```

The following statement stops a passthrough session for all remote users:

```
PASSTHROUGH STOP;
```

### Data Manipulation Language (DML)

Passthrough mode can be used to send data manipulation statements. In this case, replicated DML statements use the *before* schema before the passthrough and the *after* schema following the passthrough.

The following example drops a table on the remote database and the consolidated database.

```
-- Drop a table on the remote database
```

```
-- and at the consolidated database
PASSTHROUGH FOR Joe_Remote;
DROP TABLE CrucialData;
PASSTHROUGH STOP;
```

The following example uses PASSTHROUGH ONLY to drop a table on the remote database only.

```
-- Drop a table on the remote database only
PASSTHROUGH ONLY FOR Joe_Remote;
DROP TABLE CrucialData;
PASSTHROUGH STOP;
```

## Related Information

[Data Manipulation Statements](#)

[PASSTHROUGH Statement \[SQL Remote\]](#)

### 1.4.15 Subscription Resynchronization

When you create a remote database, you extract both the schema and data from the consolidated database and use them to build the remote database. This process ensures that each database has an initial copy of the data.

After deployment, you might consider resynchronizing subscriptions in the following circumstances:

#### **After you complete significant maintenance to the consolidated database**

For example, you make changes to the consolidated database, which updates every row in the database. By default, SQL Remote creates and sends update messages to each subscribed remote. These update messages could include the UPDATE, DELETE, and INSERT statements for each row.

If you chose to synchronize the subscription using a SYNCHRONIZE SUBSCRIPTION statement, you only send the statements required to delete all the rows in the subscribed tables and the INSERT statements to insert all new rows.

#### **When a remote database is out-of-step with the consolidated database**

If a remote database becomes out-of-step with the consolidated database, you can try to use passthrough mode.

If using passthrough mode doesn't work, you can synchronize the subscriptions. When you synchronize subscriptions, you force the remote database into step with the consolidated database. A SYNCHRONIZE SUBSCRIPTION statement includes statements to delete the contents of the subscribed tables in the remote database and statements to insert the rows of the subscription from the consolidated database to the remote database.

## Limitations

**Synchronization applies to an entire subscription**



You cannot synchronize a single table.

#### **Data loss on synchronization**

Any data on the remote database that is part of the subscription, which has not been replicated to the consolidated database, is lost.

Before synchronizing the database, use the *Unload Database Wizard* in SQL Central or the Unload utility (dbunload) to unload or back up the remote database.

#### **In this section:**

[Synchronizing Subscriptions \(SQL Central\) \[page 161\]](#)

Synchronize publications for subscribed users manually.

[Synchronizing with the SQL Remote Message Agent \(dbremote\) \[page 163\]](#)

Replace the current contents of subscribed tables with a new copy.

[Starting Subscriptions \(SQL Central\) \[page 164\]](#)

Start subscriptions to a publication for a subscribed user.

[Stopping Subscriptions \(SQL Central\) \[page 165\]](#)

Cancel a subscription for a user.

## **Related Information**

[SQL Remote Passthrough Mode \[page 157\]](#)

[Tips on Exporting Data with the Unload Database Wizard](#)

[PASSTHROUGH Statement \[SQL Remote\]](#)

[SYNCHRONIZE SUBSCRIPTION Statement \[SQL Remote\]](#)

[Unload Utility \(dbunload\)](#)

### **1.4.15.1 Synchronizing Subscriptions (SQL Central)**

Synchronize publications for subscribed users manually.

#### **Prerequisites**

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

#### **Context**

Use either the Extraction utility (dbextract) or the *Extract Database Wizard* to extract the data for the specified remote database and then manually load the data into the remote database.

### Caution

Do not run the SQL Remote Message Agent (dbremote) when running the Extraction utility (dbxtract) or the [Extract Database Wizard](#).

## Procedure

1. Shut down the SQL Remote Message Agent on the remote database and the consolidated database.
2. Connect to the consolidated database.
3. Double-click [Publications](#).
4. Double-click a publication.
5. Click the [SQL Remote Subscriptions](#) tab.
6. Manually synchronize subscriptions:
  - a. Right-click the user in the [Subscribers](#) list and click [Properties](#).
  - b. Click the [Advanced](#) tab.
  - c. Click [Synchronize Now](#).

The subscriptions are affected when you click the [Synchronize Now](#) button. Subsequently clicking [Cancel](#) on the properties window does *not* cancel the synchronize action.

7. Click [OK](#).

## Results

The subscriptions are synchronized.

## Related Information

[SYNCHRONIZE SUBSCRIPTION Statement \[SQL Remote\]](#)  
[Extraction Utility \(dbxtract\) \[page 225\]](#)

## 1.4.15.2 Synchronizing with the SQL Remote Message Agent (dbremote)

Replace the current contents of subscribed tables with a new copy.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

Use the Extraction utility (dbextract) or the *Extract Database Wizard* to synchronize subscriptions.

Extracting a large number of subscriptions, or synchronizing subscriptions to large, frequently used tables, can slow database access. You can use the SEND AT clause to specify a time to synchronize when the consolidated database is not in heavy use.

### Procedure

1. Connect to the consolidated database.
2. Execute a SYNCHRONIZE SUBSCRIPTION statement.

The SQL Remote Message Agent (dbremote) on the consolidated database sends a copy of all rows in the subscription to the subscriber. The SQL Remote Message Agent (dbremote) assumes that an appropriate database schema is in place at the remote databases.

The SQL Remote Message Agent (dbremote) on the subscriber database receives the synchronization message and it *replaces* the current contents of the subscribed tables with the new copy.

### Results

The current contents of the subscribed tables are replaced with the new copy.

### Related Information

[Send Frequency \[page 99\]](#)

[Synchronizing Subscriptions \(SQL Central\) \[page 161\]](#)

[SYNCHRONIZE SUBSCRIPTION Statement \[SQL Remote\]](#)

## 1.4.15.3 Starting Subscriptions (SQL Central)

Start subscriptions to a publication for a subscribed user.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

To start several subscriptions within a single transaction, use the REMOTE RESET statement.

### Procedure

1. In SQL Central, use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click *Publications*.
3. Double-click a publication.
4. Click the *SQL Remote Subscriptions* tab.
5. Manually synchronize subscriptions:
  - a. Right-click the user in the *Subscriber* list and click *Properties*.
  - b. Click the *Advanced* tab.
  - c. Click *Synchronize Now*.

The subscriptions are affected when you click the *Synchronize Now* button. Subsequently clicking *Cancel* on the properties window does *not* cancel the synchronize action.

### Results

The subscriptions to a publication for a subscribed user are started.

### Related Information

[START SUBSCRIPTION Statement \[SQL Remote\]](#)

[REMOTE RESET Statement \[SQL Remote\]](#)

## 1.4.15.4 Stopping Subscriptions (SQL Central)

Cancel a subscription for a user.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. Use the *SQL Anywhere 17* plug-in to connect to the database.
2. Double-click *Publications*.
3. Double-click the desired publication.
4. Click the *SQL Remote Subscriptions* tab.
5. To manually synchronize subscriptions, right-click the user in the *Subscriber* list and click *Properties*.

Click the *Advanced* tab. On this tab, click *Stop Now* to stop subscriptions.

The subscriptions are affected when you click the *Stop Now* button. Subsequently clicking *Cancel* on the properties window does *not* cancel your stop synchronization action.

### Results

The subscription is stopped.

## 1.5 Tutorial: Creating a SQL Remote System

In this tutorial you learn how to set up a SQL Remote replication system.

### Prerequisites

This tutorial refers to a user, DBA, that must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Context

In this tutorial, you complete the following steps:

1. [Lesson 1: Creating the Consolidated Database \[page 166\]](#)  
Create the consolidated database and directories for the tutorial.
2. [Lesson 2: Granting PUBLISH and REMOTE Privileges at the Consolidated Database \[page 169\]](#)  
Create the publisher for the consolidated database using Interactive SQL.
3. [Lesson 3: Creating Publications and Subscriptions \[page 170\]](#)  
Create the publication on the consolidated database using Interactive SQL.
4. [Lesson 4: Creating a SQL Remote Message Type \[page 172\]](#)  
Define the message types to use when sending data and messages for replication.
5. [Lesson 5: Extracting the Remote Database \[page 173\]](#)  
Create a database for a remote user by extracting the remote database from the consolidated database (hq).
6. [Lesson 6: Sending Data from the Consolidated Database to the Remote Database \[page 175\]](#)  
Replicate data from the consolidated database (hq) to the remote database (field) using Interactive SQL.
7. [Lesson 7: Receiving Data at the Remote Database \[page 176\]](#)  
Receive the data on the remote database (field) that was sent from the consolidated database (hq).
8. [Lesson 8: Sending Data from the Remote Database to the Consolidated Database \[page 178\]](#)  
Replicate data from the remote database (field) to the consolidated database (hq) using Interactive SQL.

### 1.5.1 Lesson 1: Creating the Consolidated Database

Create the consolidated database and directories for the tutorial.

#### Prerequisites

The DBA user must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

#### Procedure

1. Create the directories `c:\tutorial`, `c:\tutorial\hq`, and `c:\tutorial\field`.
2. From the `c:\tutorial` directory, run the following command to create the consolidated database (hq):

```
dbinit -dba DBA,passwd hq.db
```


3. Connect to the consolidated database (hq) from Interactive SQL.

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

4. Execute the following statements to create two tables in the consolidated database (hq):

```
CREATE TABLE SalesReps (  
    rep_key CHAR(12) NOT NULL,  
    name CHAR(40) NOT NULL,  
    PRIMARY KEY ( rep_key )  
);  
CREATE TABLE Customers (  
    cust_key CHAR(12) NOT NULL,  
    name CHAR(40) NOT NULL,  
    rep_key CHAR(12) NOT NULL,  
    FOREIGN KEY ( rep_key ) REFERENCES SalesReps ( rep_key ),  
    PRIMARY KEY ( cust_key )  
);
```

The following figure shows the consolidated database (hq) schema for the tutorial:



- Each sales representative is represented by one row in the SalesReps table.
- Each customer is represented by one row in the Customers table.
- Each customer is assigned to a single sales representative, and this assignment is built into the database as a foreign key from the Customers table to the SalesReps table. The relationship between the Customers table and the SalesReps table is many-to-one.

Table name	Description
<a href="#">SalesReps</a>	In the SalesReps table, there is a row for each sales representative that works for the company. The SalesReps table has the following columns:  <b>rep_key</b> An identifier for each sales representative. This is the primary key.  <b>name</b> The name of each sales representative.

Table name	Description
<i>Customers</i>	<p>In the Customers table, there is one row for each customer that does business with the company. The Customers table includes the following columns:</p> <p><b>cust_key</b> An identifier for each customer. This is the primary key.</p> <p><b>name</b> The name of each customer.</p> <p><b>rep_key</b> An identifier for the sales representative in a sales relationship. This is a foreign key to the SalesReps table.</p>

5. Execute the following statements to add sample data to the SalesReps and Customers tables:

```
INSERT INTO SalesReps ( rep_key, name ) VALUES ( 'rep1', 'Field User' );
INSERT INTO SalesReps ( rep_key, name ) VALUES ( 'rep2', 'Another User' );
INSERT INTO Customers ( cust_key, name, rep_key ) VALUES ( 'cust1', 'Ocean Sports', 'rep1' );
INSERT INTO Customers ( cust_key, name, rep_key ) VALUES ( 'cust2', 'Sports Plus', 'rep2' );
COMMIT;
```

6. Execute the following statements to confirm that the tables were created:

```
SELECT * FROM SalesReps;
```

The above query returns the following data from the SalesReps table:

rep_key	name
rep1	Field User
rep2	Another User

```
SELECT * FROM Customers;
```

The above query returns the following data from the Customers table:

cust_key	name	rep_key
cust1	Ocean Sports	rep1
cust2	Sports Plus	rep2

## Results

The tables are created and populated with data.



## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Creating a SQL Remote System \[page 165\]](#)

**Next task:** [Lesson 2: Granting PUBLISH and REMOTE Privileges at the Consolidated Database \[page 169\]](#)

## 1.5.2 Lesson 2: Granting PUBLISH and REMOTE Privileges at the Consolidated Database

Create the publisher for the consolidated database using Interactive SQL.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

Every database in a SQL Remote system requires a publisher, which is a unique user with PUBLISH privilege. All outgoing SQL Remote messages, including publication updates and receipt confirmations, are identified by their publisher. Every database in a SQL Remote system sends receipt confirmations.

### Procedure

1. If you are not currently connected to the consolidated database (hq), run the following command:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

2. Execute the following statements to create the user hq\_user that has PUBLISH privilege:

```
CREATE USER hq_user;  
GRANT PUBLISH TO hq_user;
```

3. Execute the following statement to check the publishing user ID of the database, which should be hq\_user:

```
SELECT CURRENT PUBLISHER;
```

4. A database, such as a consolidated database, that sends messages to other databases must specify which remote databases it sends messages to. To specify these remote databases on the consolidated database, grant REMOTE privilege to the publishers of the remote databases. REMOTE privilege identifies databases

that receive messages from the current database. Execute the following statements to create the remote user `field_user` that has `REMOTE` privilege:

```
CREATE USER field_user;  
GRANT REMOTE TO field_user TYPE FILE ADDRESS 'field';
```

For a `FILE` message type, the `ADDRESS` value must be a valid folder path. For an `SMTP` message type, the `ADDRESS` value must be a valid email address.

## Results

The users that identify the publisher of the consolidated database and the remote database are created.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Creating a SQL Remote System \[page 165\]](#)

**Previous task:** [Lesson 1: Creating the Consolidated Database \[page 166\]](#)

**Next task:** [Lesson 3: Creating Publications and Subscriptions \[page 170\]](#)

## 1.5.3 Lesson 3: Creating Publications and Subscriptions

Create the publication on the consolidated database using Interactive SQL.

### Prerequisites

You must have the `SYS_REPLICATION_ADMIN_ROLE` system role.

### Context

A publication describes the set of data to be replicated. In this lesson, you create a publication named `SalesRepData` that replicates all rows of the `SalesReps` table and some rows of the `Customers` table. You subscribe a user to a publication by creating a subscription.

## Procedure

1. If you are not currently connected to the consolidated database (hq), then run the following command:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

2. Execute the following statement to create a publication named SalesRepData:

```
CREATE PUBLICATION SalesRepData (  
    TABLE SalesReps,  
    TABLE Customers SUBSCRIBE BY rep_key  
);
```

The SalesRepData publication publishes:

- The entire SalesReps table
  - All of the columns in the Customers table but only the rows that match a specified rep\_key value
3. Execute the following statement to create a subscription to SalesRepData for field\_user:

```
CREATE SUBSCRIPTION TO SalesRepData ( 'rep1' ) FOR field_user;
```

The value rep1 is the rep\_key value for field\_user in the SalesReps table.

### i Note

In this tutorial, we do not discuss how to protect against duplicate entries of primary key values, one arising in the consolidated database and the other in the remote database. In general, this is resolved through the use of primary key pools. Each remote is assigned a unique pool of keys to draw from so that newly inserted rows have keys that do not collide with the consolidated or other remotes. One way of creating these pools is through the use of the "SET OPTION public.global\_database\_id=value" statement.

## Results

The SalesRepData publication is created to replicate all rows of the SalesReps table, and some rows of the Customers table to field\_user.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Creating a SQL Remote System \[page 165\]](#)

**Previous task:** [Lesson 2: Granting PUBLISH and REMOTE Privileges at the Consolidated Database \[page 169\]](#)

**Next task:** [Lesson 4: Creating a SQL Remote Message Type \[page 172\]](#)

## Related Information

[Create a SQL Remote System \[page 11\]](#)

## 1.5.4 Lesson 4: Creating a SQL Remote Message Type

Define the message types to use when sending data and messages for replication.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

A message type description has two parts:

#### A message system supported by SQL Remote

This tutorial uses the FILE message system. The FILE message system is a simple file-sharing system

#### A FILE address

A user's FILE address is a subdirectory to which all their incoming messages are sent. An application retrieves the messages from this directory. In this tutorial the FILE address of the consolidated database is hq and it is a subdirectory of c:\tutorial.

### Procedure

1. If you are not currently connected to the consolidated database (hq), then run the following command:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

2. Execute the following statement to create a FILE message type:

```
CREATE REMOTE MESSAGE TYPE FILE ADDRESS 'hq';
```

For a FILE message type, the ADDRESS value must be a valid folder path. For an SMTP message type, the ADDRESS value must be a valid email address.

## Results

The FILE message type is created.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Creating a SQL Remote System \[page 165\]](#)

**Previous task:** [Lesson 3: Creating Publications and Subscriptions \[page 170\]](#)

**Next task:** [Lesson 5: Extracting the Remote Database \[page 173\]](#)

## 1.5.5 Lesson 5: Extracting the Remote Database

Create a database for a remote user by extracting the remote database from the consolidated database (hq).

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

You must configure the remote database to send and receive messages and participate in a SQL Remote system. Like the consolidated database (hq), the remote database needs a CURRENT PUBLISHER to identify the source of outgoing messages. It also needs to have the consolidated database (hq) identified as a subscriber.

Run the dbxtract utility to create a remote database that contains:

- a subscription to the consolidated database
- a publication
- a current copy of the data

## Procedure

1. Extract the remote database schema and data from the consolidated database (hq) for the user `field_user` by running the following command from the `c:\tutorial` directory:

```
dbxtract -v -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=C:\tutorial\hq.db" -  
an c:\tutorial\field.db field_user
```

This command:

- Starts the subscription to the remote user in the consolidated database.
- Creates a new, empty database at `c:\tutorial\field.db` and performs the following tasks:
  - Creates a message type at the remote database (field).
  - Grants PUBLISH privilege to the remote database (field).
  - Creates the SalesReps and Customers tables in the remote database (field). These tables contain the same data as the consolidated database (hq).
  - Creates a publication to identify the data being replicated at the remote database (field).
  - Creates the subscription for the consolidated database (hq) and starts the subscription.

### ⚠ Caution

In a production environment, do not store two replicating databases in the same directory. The consolidated database should not be stored with the remote database in case of media failure.

2. Connect to the newly created remote database with the following command:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=field_db;DBF=c:\tutorial\field.db"
```

3. Execute the following statements to confirm that the tables were created:

```
SELECT * FROM SalesReps;
```

The above query returns the following data from the SalesReps table:

rep_key	name
rep1	Field User
rep2	Another User

```
SELECT * FROM Customers;
```

The above query returns the following data from the Customers table:

cust_key	name	rep_key
cust1	Ocean Sports	rep1

## Results

The remote database is created for the remote user.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Creating a SQL Remote System \[page 165\]](#)

**Previous task:** [Lesson 4: Creating a SQL Remote Message Type \[page 172\]](#)

**Next task:** [Lesson 6: Sending Data from the Consolidated Database to the Remote Database \[page 175\]](#)

## 1.5.6 Lesson 6: Sending Data from the Consolidated Database to the Remote Database

Replicate data from the consolidated database (hq) to the remote database (field) using Interactive SQL.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

Many.

### Procedure

1. If you are not currently connected to the consolidated database (hq), then run the following command:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

2. Execute the following statements to add sample data to the SalesReps and Customers tables:

```
INSERT INTO SalesReps ( rep_key, name ) VALUES ( 'rep3', 'Example User' );
INSERT INTO Customers ( cust_key, name, rep_key ) VALUES ( 'cust3', 'Land
Sports', 'repl' );
INSERT INTO Customers ( cust_key, name, rep_key ) VALUES ( 'cust4', 'Air
Plus', 'rep2' );
COMMIT;
```

3. Execute the following statements to confirm that the data has been entered:

```
SELECT * FROM SalesReps;
```

```
SELECT * FROM Customers;
```

4. To send the rows to the remote database (field), run the Message Agent on the consolidated database (hq) from the `c:\tutorial` directory:

```
dbremote -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

5. When the Message Agent window displays `Execution Completed`, click *Shutdown*.
6. Browse to `c:\tutorial\field`.

A file named `hq.0` is listed in the directory. This file contains the changes sent from the consolidated database (hq).

## Results

The sample data is added to the `SalesReps` and `Customers` tables and sent from the consolidated database to the remote database.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Creating a SQL Remote System \[page 165\]](#)

**Previous task:** [Lesson 5: Extracting the Remote Database \[page 173\]](#)

**Next task:** [Lesson 7: Receiving Data at the Remote Database \[page 176\]](#)

## 1.5.7 Lesson 7: Receiving Data at the Remote Database

Receive the data on the remote database (field) that was sent from the consolidated database (hq).

### Prerequisites

You must have the `SYS_REPLICATION_ADMIN_ROLE` system role.



## Procedure

1. If you are not currently connected to the remote database (field), then run the following command:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=field_db;DBF=c:\tutorial\field.db"
```

2. At the remote database (field), run the Message Agent from the c:\tutorial directory:

```
dbremote -c "UID=DBA;PWD=passwd;SERVER=field_db;DBF=c:\tutorial\field.db"
```

3. When the Message Agent window displays Execution Completed, click [Shutdown](#).

The c:\tutorial\field\hq.0 file has been replaced by a file named c:\tutorial\hq\field.0. The field.0 file contains the receipt confirmation.

4. Verify that the remote database (field) contains data.
  - a. Execute the following statement to view the contents of the SalesReps table:

```
SELECT * FROM SalesReps;
```

The SalesReps table contains all rows entered at the consolidated database (hq). This is because the SalesRepData publication included all the data from the SalesReps table.

rep_key	name
rep1	Field User
rep2	Another User
rep3	Example User

- b. Execute the following statement to view the contents of the Customers table:

```
SELECT * FROM Customers;
```

The Customers table now also contains a row with the Land Sports customer data that was entered at the consolidated database (hq).

cust_key	name	rep_key
cust1	Ocean Sports	rep1
cust3	Land Sports	rep1

5. At the consolidated database (hq) run the Message Agent from the c:\tutorial directory:

```
dbremote -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

In the c:\tutorial\hq directory, the file field.0 disappears.

## Results

The data that was sent from the consolidated database is received on the remote database.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Creating a SQL Remote System \[page 165\]](#)

**Previous task:** [Lesson 6: Sending Data from the Consolidated Database to the Remote Database \[page 175\]](#)

**Next task:** [Lesson 8: Sending Data from the Remote Database to the Consolidated Database \[page 178\]](#)

## 1.5.8 Lesson 8: Sending Data from the Remote Database to the Consolidated Database

Replicate data from the remote database (field) to the consolidated database (hq) using Interactive SQL.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. If you are not currently connected to the remote database (field), then run the following command:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=field_db;DBF=c:\tutorial\field.db"
```

2. Execute the following statement to insert a row at the remote database (field):

```
INSERT INTO Customers ( cust_key, name, rep_key ) VALUES ( 'cust5', 'North  
Land Trading', 'repl' );  
COMMIT;
```

3. From the c:\tutorial directory, run the dbremote utility against the remote database (field):

```
dbremote -c "UID=DBA;PWD=passwd;SERVER=field_db;DBF=c:\tutorial\field.db"
```

In the c:\tutorial\hq directory, the file field.1 appears.

4. If you are not currently connected to the consolidated database (hq), then run the following command:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

5. At the consolidated database (hq) run the Message Agent from the c:\tutorial directory:

```
dbremote -c "UID=DBA;PWD=passwd;SERVER=server_hq;DBF=c:\tutorial\hq.db"
```

- When the Message Agent window displays `Execution Completed`, click [Shutdown](#).
- Browse to `c:\tutorial\field`.

The `hq.1` file has been replaced by a file named `hq.2`. The `hq.2` file contains the receipt confirmation.

- Execute the following statement to view the data in the `Customers` table in the consolidated database (`hq`):

```
SELECT * FROM Customers;
```

This query returns the following results:

<code>cust_key</code>	<code>name</code>	<code>rep_key</code>
<code>cust1</code>	Ocean Sports	<code>rep1</code>
<code>cust2</code>	Sports Plus	<code>rep2</code>
<code>cust3</code>	Land Sports	<code>rep1</code>
<code>cust4</code>	Air Plus	<code>rep2</code>
<code>cust5</code>	North Landing Trading	<code>rep1</code>

## Results

The data is replicated from the remote database to the consolidated database.

**Task overview:** [Tutorial: Creating a SQL Remote System \[page 165\]](#)

**Previous task:** [Lesson 7: Receiving Data at the Remote Database \[page 176\]](#)

## 1.6 Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server

In this tutorial you learn how to set up a SQL Remote replication system that uses the HTTP messaging system..

### Prerequisites

This tutorial refers to a user, `DBA`, that must have the `SYS_REPLICATION_ADMIN_ROLE` system role.

## Context

The consolidated database uses the FILE message system to replicate changes, while the remote database uses the HTTP message system to replicate changes.

In this tutorial, you:

- Create a consolidated SQL Anywhere database and a remote SQL Anywhere database that contains all the data in the consolidated database.
  - Create a directory structure to store the messages generated by SQL Remote. The consolidated database accesses the files using the FILE message system, while the remote database uses the HTTP message system.
  - Configure the consolidated database to act as the message server for the HTTP message system.
  - Create a remote database that sends messages using the HTTP messaging system.
  - Replicate data between the consolidated and remote databases.
1. [Lesson 1: Creating the Consolidated Database \[page 180\]](#)  
Create the directories needed to store the databases and their transactions logs, as well as the directory structure for the messages. You also define the schema of the consolidated database, including creation of the remote user and the publication and subscription needed to replicate data.
  2. [Lesson 2: Configuring the Consolidated Database to Act as the Message Server \[page 183\]](#)  
Configure the consolidated database to act as the message server for the HTTP message system.
  3. [Lesson 3: Creating the Remote Database \[page 184\]](#)  
Extract the remote database, and then replace the FILE message system at the remote database with the HTTP message system.
  4. [Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases \[page 186\]](#)  
Add data to the consolidated and remote database, run SQL Remote to replicate the changes, and then confirm that the data is consistent in both databases.
  5. [Lesson 5: Finishing up \[page 189\]](#)  
Shut down the remote and consolidated databases.

### 1.6.1 Lesson 1: Creating the Consolidated Database

Create the directories needed to store the databases and their transactions logs, as well as the directory structure for the messages. You also define the schema of the consolidated database, including creation of the remote user and the publication and subscription needed to replicate data.

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Context

When SQL Remote runs against the consolidated database, it uses the FILE message system to send and receive messages, but the remote database uses the HTTP message system.

### i Note

For secure data transfer, use the HTTPS protocol. HTTPS requires certificates, so this tutorial uses HTTP to simplify the steps.

## Procedure

1. Create the following directories to hold the consolidated database and the remote database:
  - `c:\tutorial`
  - `c:\tutorial\consol`
  - `c:\tutorial\remote`
2. Create the following directories to hold the message files generated by the consolidated and remote databases:
  - `c:\tutorial\messages`
  - `c:\tutorial\messages\conmsgs`
  - `c:\tutorial\messages\remmsgs`
3. In the `c:\tutorial\consol` directory, run the following command to create the consolidated database:

```
dbinit -dba DBA,passwd consoldb.db
```

4. Start the consolidated database:

```
dbeng17 consoldb.db -xs http(port=8033)
```

`-xs http(8033)` is required on the command line because this is the database server that will be accepting HTTP requests from the remote database and accessing the message files that exist under the `c:\tutorial\messages` directory. No web services are defined at this time. In this lesson, you only start the personal database server, so only SQL Remote processes on this computer are able to communicate with the message server using HTTP. In a production environment, you would typically use the network server so that SQL Remote processes on other computers would also have access to the web services.

5. Using Interactive SQL, connect to the consolidated database as a user with the `SYS_REPLICATION_ADMIN_ROLE` system role:

```
dbisql -c "SERVER=consoldb;UID=DBA;PWD=passwd"
```

6. To set the global database ID for the consolidated database, execute the following statement (the global database ID is needed so that distinct primary keys are chosen for all databases when using the GLOBAL AUTOINCREMENT default):

```
SET OPTION public.global_database_id=0;
```

- The schema for the database in this tutorial consists of a single table and all the columns and rows from the table replicate to every remote user. Execute the following statements for the consolidated database to create the single table in the database:

```
CREATE TABLE employees (
    employee_id BIGINT NOT NULL DEFAULT GLOBAL AUTOINCREMENT(1000000)
    PRIMARY KEY,
    first_name VARCHAR(128) NOT NULL,
    last_name VARCHAR(128) NOT NULL,
    hire_date TIMESTAMP NOT NULL DEFAULT TIMESTAMP
);
```

- Execute the following statements on the consolidated database to add sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Kelly', 'Meloy');
INSERT INTO employees (first_name, last_name) VALUES ('Melisa', 'Boysen');
COMMIT;
```

- Execute the following statement on the consolidated database to confirm that the table was created and populated with data:

```
SELECT * FROM employees;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the values you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	2018-07-06 10:30:02.369
2	Melisa	Boysen	2018-07-06 10:30:02.371

- Execute the following statements to create a user without a password that has the PUBLISH privilege:

```
CREATE USER ConUser;
GRANT PUBLISH TO ConUser;
```

- For performance reasons, the HTTP message system is only used at the remote database, and not at the consolidated. The following statements configure the use of the FILE-based message system at the consolidated database:

```
CREATE REMOTE MESSAGE TYPE FILE ADDRESS 'conmsgs';
SET REMOTE FILE OPTION public.directory='c:\\tutorial\\messages';
SET REMOTE FILE OPTION public.debug='yes';
```

- Execute the following statements to create a remote user with a password that has the REMOTE privilege and uses the FILE message system with the specified file path (that is located under the directory specified by the public.directory REMOTE FILE OPTION setting):

```
CREATE USER RemUser IDENTIFIED BY passwd;
GRANT REMOTE TO RemUser TYPE FILE ADDRESS 'remmsgs';
```

- A publication describes the set of data to be replicated. Create a publication named pub\_employees that replicates all rows of the employees table. You subscribe a user to a publication by creating a subscription.

```
CREATE PUBLICATION pub_employees ( TABLE employees );
CREATE SUBSCRIPTION TO pub_employees FOR RemUser;
```

## Results

The directories needed to store the databases and their transaction logs are created, as well as the directory structure for the messages. The schema of the consolidated database is defined, including creation of the remote user and the publication and subscription needed to replicate data.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server \[page 179\]](#)

**Next task:** [Lesson 2: Configuring the Consolidated Database to Act as the Message Server \[page 183\]](#)

## 1.6.2 Lesson 2: Configuring the Consolidated Database to Act as the Message Server

Configure the consolidated database to act as the message server for the HTTP message system.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

It is also possible to configure a separate database and database server to act as the message server.

### Procedure

1. If you have not already done so, connect to the consolidated database using Interactive SQL as a user with the SYS\_REPLICATION\_ADMIN\_ROLE system role:

```
dbisql -c "SERVER=consolddb;UID=DBA;PWD=passwd"
```

2. When a database is first initialized, none of the web services needed to accept HTTP requests from remote users is defined, and neither are definitions to allow the database server to access the directory where the

message files are stored. The creation of these objects is automated with the use of the `sr_add_message_server` stored procedure, which takes an optional parameter to specify who owns all the objects. Execute the following statements on the consolidated database to define all the objects needed for the message server, and specify that all the objects are owned by the consolidated database user `ConUser`:

```
CREATE ROLE FOR USER ConUser;  
SET REMOTE HTTP OPTION ConUser.root_directory='c:\\tutorial\\messages';  
CALL sr_add_message_server( 'ConUser' );  
COMMIT;
```

## Results

The consolidated database is configured to act as the message server for the HTTP message system.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server \[page 179\]](#)

**Previous task:** [Lesson 1: Creating the Consolidated Database \[page 180\]](#)

**Next task:** [Lesson 3: Creating the Remote Database \[page 184\]](#)

## 1.6.3 Lesson 3: Creating the Remote Database

Extract the remote database, and then replace the FILE message system at the remote database with the HTTP message system.

## Prerequisites

You must have the `SYS_REPLICATION_ADMIN_ROLE` system role.



## Procedure

1. In the `c:\tutorial\remote` directory, run the following command to create the remote database:

```
dbinit -dba DBA,passwd remotedb.db
```

2. Start the remote database:

```
dbeng17 remotedb.db
```

Ensure that the consolidated database is also running (it should be).

3. Use `dbxtract` to set up the remote database. Run the following command to extract database objects for the remote user `RemUser` from the consolidated database:

```
dbxtract -xx -ac "SERVER=remotedb;UID=DBA;PWD=passwd" -c  
"SERVER=consolddb;UID=DBA;PWD=passwd" RemUser
```

4. Connect to the remote database using the following command:

```
dbisql -c "SERVER=remotedb;UID=DBA;PWD=passwd"
```

5. The consolidated database uses the FILE message system, so when `dbxtract` runs, it creates SQL Remote definitions assuming that the remote database also uses the FILE message system. Since the remote database will use the HTTP message system, execute the following statements on the remote database to reset the FILE message system for this remote database:

```
CREATE REMOTE MESSAGE TYPE FILE ADDRESS '';  
SET REMOTE FILE OPTION public.directory='';  
SET REMOTE FILE OPTION public.debug='';
```

6. Execute the following statements on the remote database to configure the HTTP message system for this remote database:

```
CREATE REMOTE MESSAGE TYPE HTTP ADDRESS 'remmsgs';  
GRANT CONSOLIDATE TO ConUser TYPE HTTP ADDRESS 'conmsgs';  
SET REMOTE HTTP OPTION public.user_name='RemUser';  
SET REMOTE HTTP OPTION public.password='passwd';  
SET REMOTE HTTP OPTION public.debug='yes';  
SET REMOTE HTTP OPTION public.https='no';  
SET REMOTE HTTP OPTION public.url='localhost:8033';  
COMMIT;
```

7. Verify that the `employees` table in the remote database contains the two rows of data that existed in the consolidated database after the extraction. Execute the following statement to view the contents of the `employees` table:

```
SELECT * FROM employees;
```

The query returns the following data from the `employees` table, although the `hire_date` column contains the time you inserted the row, and not the data you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	2018-07-06 10:30:02.369
2	Melisa	Boysen	2018-07-06 10:30:02.371

## Results

The remote database is extracted and the FILE message system at the remote database is replaced with the HTTP message system.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server \[page 179\]](#)

**Previous task:** [Lesson 2: Configuring the Consolidated Database to Act as the Message Server \[page 183\]](#)

**Next task:** [Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases \[page 186\]](#)

## 1.6.4 Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases

Add data to the consolidated and remote database, run SQL Remote to replicate the changes, and then confirm that the data is consistent in both databases.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. If you are not currently connected to the consolidated database, run the following command:

```
dbisql -c "SERVER=consolddb;UID=DBA;PWD=passwd"
```

2. Execute the following statements on the consolidated database to add additional sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Javier', 'Sporr');  
COMMIT;
```

3. If you are not currently connected to the remote database, then run the following command:

```
dbisql -c "SERVER=remotedb;UID=DBA;PWD=passwd"
```

4. Execute the following statements on the remote database to add additional sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Nelson', 'Kreitzer');  
COMMIT;
```

5. For the consolidated database, run the Message Agent:

```
dbremote -c "SERVER=consolddb;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial  
\conlog1.txt
```

This command initiates a scan of the consolidated database transaction log and generates a message for the remote database using the FILE message system. Since the debug message system parameter has been set for the FILE message system in the consolidated database, you can look at the `c:\tutorial\conlog1.txt` file and verify that you see debug messages indicating that messages are being written to the `c:\tutorial\messages\remmsgs` directory. For example:

```
I. 2018-07-06 15:19:27. Processing transactions from active transaction log  
I. 2018-07-06 15:19:27. Sending message to "RemUser"  
(0-0000000000-0000720309-0)  
I. 2018-07-06 15:19:27. sopen "c:\tutorial\messages\remmsgs\conmsgs.0"  
I. 2018-07-06 15:19:27. write "c:\tutorial\messages\remmsgs\conmsgs.0"  
I. 2018-07-06 15:19:27. close "c:\tutorial\messages\remmsgs\conmsgs.0"
```

6. For the remote database, run the Message Agent:

```
dbremote -c "SERVER=remotedb;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial  
\remlog.txt
```

Using the HTTP message system, this command receives and applies the message that was just generated by the consolidated database. It then scans the transaction log and sends a message back to the consolidated database with the new row that was added in the remote database. Since the debug message system parameter has been set for the HTTP message system in the remote database, you can look at the `c:\tutorial\remlog.txt` file and verify that you see debug messages indicating that the HTTP message system is being used. For example:

```
I. 2018-07-06 15:24:36. Sending message to "ConUser"  
(0-0000000000-0000727798-0)  
I. 2018-07-06 15:24:36. HTTPWriteMessage "remmsgs.0"  
I. 2018-07-06 15:24:36. HTTPWriteMessage: success -- filename "remmsgs.0"  
I. 2018-07-06 15:24:37. HTTPDisconnect
```

7. For the consolidated database, run the Message Agent again specifying a different log file:

```
dbremote -c "SERVER=consolddb;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial  
\conlog2.txt
```

This command receives and applies the message that was just generated by the remote database using the FILE-based message system.

- To verify that the consolidated database contains all four rows of data, execute the following statement to view the contents of the employees table:

```
SELECT * FROM employees
ORDER BY employee_id;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the values you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	2018-07-06 10:30:02.369
2	Melisa	Boysen	2018-07-06 10:30:02.371
3	Javier	Spoor	2018-07-06 13:52:03.502
102000001	Nelson	Kreitzer	2018-07-06 13:52:14.063

- Disconnect from the consolidated database.
- Verify that the remote database contains all four rows of data by executing the following statement to view the contents of the employees table:

```
SELECT * FROM employees
ORDER BY employee_id;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the data you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	2018-07-06 10:30:02.369
2	Melisa	Boysen	2018-07-06 10:30:02.371
3	Javier	Spoor	2018-07-06 13:52:03.502
102000001	Nelson	Kreitzer	2018-07-06 13:52:14.063

- Disconnect from the remote database.

## Results

Data is added to the consolidated and remote database, the changes are replicated, and the data consistency has been verified.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server \[page 179\]](#)

**Previous task:** [Lesson 3: Creating the Remote Database \[page 184\]](#)

**Next task:** [Lesson 5: Finishing up \[page 189\]](#)

## 1.6.5 Lesson 5: Finishing up

Shut down the remote and consolidated databases.

### Prerequisites

You must have the SERVER OPERATOR system privilege to shut down a network server. No privilege is required to shut down a personal server.

### Procedure

1. Run the following command to shut down the remote database:

```
dbstop -y -c "SERVER=remotedb;UID=DBA;PWD=passwd"
```

2. Run the following command to shut down the consolidated database:

```
dbstop -y -c "SERVER=consoldb;UID=DBA;PWD=passwd"
```

### Results

The remote and consolidated databases are shut down.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with the Consolidated Database as the Message Server \[page 179\]](#)

**Previous task:** [Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases \[page 186\]](#)

## 1.7 Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server

In this tutorial you learn how to set up a SQL Remote replication system that uses the HTTP messaging system with a separate message server.

### Prerequisites

This tutorial refers to a user, DBA, that must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

The consolidated database uses the FILE message system to replicate changes, while the remote database uses the HTTP message system to replicate changes.

In this tutorial, you:

- Create a consolidated SQL Anywhere database and a remote SQL Anywhere database that contains all the data in the consolidated database.
- Create a directory structure to store the messages generated by SQL Remote. The consolidated database accesses the files using the FILE message system, while the remote database uses the HTTP message system.
- Create a message server SQL Anywhere database to act as a web server to receive messages from the remote database using the HTTP protocol.
- Create a remote database that sends messages using the HTTP messaging system
- Replicate data between the consolidated and remote databases.

1. [Lesson 1: Creating the Consolidated Database \[page 191\]](#)

Create the directories needed to store the databases and their transactions logs, as well as the directory structure for the messages. You also define the schema of the consolidated database, including creation of the remote user and the publication and subscription needed to replicate data.

2. [Lesson 2: Creating the Message Server \[page 193\]](#)

Use a separate database server to host message servers. This helps distribute the amount of work performed to process messages between the two database servers, and also adds a level of security because you haven't opened up HTTP access to your consolidated database.

3. [Lesson 3: Creating the Remote Database \[page 196\]](#)

Extract the remote database and replace the FILE message system at the remote database with the HTTP message system.

4. [Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases \[page 198\]](#)

Add data to the consolidated and remote database, run SQL Remote to replicate the changes, and then confirm that the data is consistent in both databases.

5. [Lesson 5: Finishing up \[page 200\]](#)

Shut down the three database servers you started in this tutorial.

## 1.7.1 Lesson 1: Creating the Consolidated Database

Create the directories needed to store the databases and their transactions logs, as well as the directory structure for the messages. You also define the schema of the consolidated database, including creation of the remote user and the publication and subscription needed to replicate data.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

When SQL Remote runs against the consolidated database, it uses the FILE message system to send and receive messages, but the remote database uses the HTTP message system.

### Procedure

1. Create the following directories to hold the consolidated database, the remote database, and the message server database:
  - c:\tutorial
  - c:\tutorial\consol
  - c:\tutorial\remote
  - c:\tutorial\msgsrv
2. Create the following directories to hold the message files generated by the consolidated database and the remote database:
  - c:\tutorial\messages
  - c:\tutorial\messages\conmsgs
  - c:\tutorial\messages\remmsgs
3. In the c:\tutorial\consol directory, run the following command to create the consolidated database:

```
dbinit -dba DBA,passwd consoldb.db
```

4. Start the consolidated database:

```
dbeng17 consoldb.db
```

- Using Interactive SQL, connect to the consolidated database as a user with the SYS\_REPLICATION\_ADMIN\_ROLE system role:

```
dbisql -c "SERVER=consolddb;UID=DBA;PWD=passwd"
```

- To set the global database ID for the consolidated database, execute the following statement (the global database ID is needed so that distinct primary keys are chosen for all databases when using the GLOBAL AUTOINCREMENT default):

```
SET OPTION public.global_database_id=0;
```

- The schema for the database in this tutorial consists of a single table and all the columns and rows from the table replicate to every remote user. Execute the following statements on the consolidated database to create the single table in the database:

```
CREATE TABLE employees (  
    employee_id BIGINT NOT NULL DEFAULT GLOBAL AUTOINCREMENT(1000000)  
    PRIMARY KEY,  
    first_name VARCHAR(128) NOT NULL,  
    last_name VARCHAR(128) NOT NULL,  
    hire_date TIMESTAMP NOT NULL DEFAULT TIMESTAMP  
);
```

- Execute the following statements on the consolidated database to add sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Kelly', 'Meloy');  
INSERT INTO employees (first_name, last_name) VALUES ('Melisa', 'Boysen');  
COMMIT;
```

- Execute the following statement on the consolidated database to confirm that the table was created and populated with data:

```
SELECT * FROM employees;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the values you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	2018-07-06 10:30:02.369
2	Melisa	Boysen	2018-07-06 10:30:02.371

- In this tutorial, the publisher and remote users are not assigned passwords, so while the users exist in the database, you cannot connect to the database with these users. Execute the following statements to create a user that has PUBLISH privileges:

```
CREATE USER ConUser;  
GRANT PUBLISH TO ConUser;
```

- For performance reasons, the HTTP message system can only be used at the remote database, and not at the consolidated. The following statements configure the use of the FILE-based message system at the consolidated database:

```
CREATE REMOTE MESSAGE TYPE FILE ADDRESS 'conmsgs';  
SET REMOTE FILE OPTION public.directory='c:\\tutorial\\messages';  
SET REMOTE FILE OPTION public.debug='yes';
```



- Execute the following statements to create a remote user without a password that has the REMOTE privilege and uses the FILE message system with the specified file path:

```
CREATE USER RemUser;  
GRANT REMOTE TO RemUser TYPE FILE ADDRESS 'remmsgs';
```

- A publication describes the set of data to be replicated. Create a publication named pub\_employees that replicates all rows of the employees table. You subscribe a user to a publication by creating a subscription.

```
CREATE PUBLICATION pub_employees ( TABLE employees );  
CREATE SUBSCRIPTION TO pub_employees FOR RemUser;
```

## Results

The directories needed to store the databases and their transactions logs are created, as well as the directory structure for the messages. The schema of the consolidated database is defined, including creation of the remote user and the publication and subscription needed to replicate data.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server \[page 190\]](#)

**Next task:** [Lesson 2: Creating the Message Server \[page 193\]](#)

## 1.7.2 Lesson 2: Creating the Message Server

Use a separate database server to host message servers. This helps distribute the amount of work performed to process messages between the two database servers, and also adds a level of security because you haven't opened up HTTP access to your consolidated database.

## Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Context

### i Note

For secure data transfer, use the HTTPS protocol. HTTPS requires certificates, so this tutorial uses HTTP to simplify the steps.

## Procedure

1. In the `c:\tutorial\msgsrv` directory, run the following command to create the message server database:

```
dbinit -dba DBA,passwd msgsrvdb.db
```

2. Start the message server database:

```
dbeng17 msgsrvdb.db -xs http(port=8033)
```

`-xs http(8033)` is required on the command line because this is the database server that accepts HTTP requests from the remote database and accesses the messages files that exist under the `c:\tutorial\messages` directory. While no web services have been defined at the time the database server starts, they are created in this lesson. As well, only the personal database server has been started, so only SQL Remote processes on this computer can communicate with the message server using HTTP. In a production environment, you would typically use the network server so that SQL Remote processes on other computers would also have access to the web services.

3. When you create a separate message server, you need to copy much of the schema of the consolidated database into the message server database, particularly information about the remote users that are defined and their addresses. While you can do this manually, the easiest way to accomplish this task is to use the `dbunload` utility to set up message server database with the same schema as the consolidated database:

```
dbunload -n -xx -ac "SERVER=msgsrvdb;UID=DBA;PWD=passwd" -c  
"SERVER=consoldb;UID=DBA;PWD=passwd"
```

The options used in the `dbunload` command do the following:

#### **-n**

Indicates that only the schema is to be unloaded, and none of the data in the consolidated database is added to the message server database.

#### **-xx**

Performs an external unload and reload, which is needed when both databases involved are already running.

#### **-ac "SERVER=msgsrvdb;UID=DBA;PWD=passwd"**

Specifies the destination connection for the unload, which for this lesson is the message server database.

#### **-c "SERVER=cons;DBN=cons;UID=DBA;PWD=passwd"**

Specifies the source connection for the unload, which for this lesson is the consolidated database.

- Using Interactive SQL, connect to the message server database as a user with the SYS\_REPLICATION\_ADMIN\_ROLE system role:

```
dbisql -c "SERVER=msgsrvdb;UID=DBA;PWD=passwd"
```

Because the publisher ConUser and remote user RemUser do not yet have passwords, neither of those users is able to connect to the consolidated database. A password is required for these users in the message server database, since the HTTP requests that come from remote users use the publisher of the remote database and the password that you supply to authenticate with the message server database. Execute the following statements on the message server database to define passwords for the publisher and remote user:

```
ALTER USER ConUser IDENTIFIED BY passwd;  
ALTER USER RemUser IDENTIFIED BY passwd;
```

- When a database is first initialized, none of the web services needed to accept HTTP requests from remote users is defined, and neither are definitions to allow the database server to access the directory where the message files are stored. The creation of these objects is automated with the use of the sr\_add\_message\_server stored procedure, which takes an optional parameter to specify who owns all the objects. Execute the following statements for the message server database to define all the objects needed for the message server and specify that all the objects are owned by the ConUser user:

```
CREATE ROLE FOR USER ConUser;  
SET REMOTE HTTP OPTION ConUser.root_directory='c:\\tutorial\\messages';  
CALL sr_add_message_server( 'ConUser' );  
COMMIT;
```

## Results

A separate database server now hosts the message server database.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server \[page 190\]](#)

**Previous task:** [Lesson 1: Creating the Consolidated Database \[page 191\]](#)

**Next task:** [Lesson 3: Creating the Remote Database \[page 196\]](#)

## Related Information

[sr\\_add\\_message\\_server System Procedure \[page 239\]](#)

[-xs Database Server Option](#)

## 1.7.3 Lesson 3: Creating the Remote Database

Extract the remote database and replace the FILE message system at the remote database with the HTTP message system.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. In the `c:\tutorial\remote` directory, run the following command to create the remote database:

```
dbinit -dba DBA,passwd remotedb.db
```

2. Start the remote database:

```
dbeng17 remotedb.db
```

Ensure that the consolidated database is also running (it should be).

3. Use `dbxtract` to create the remote database. Run the following command to extract the database objects for the `RemUser` user from the consolidated database, and leave the database server for the remote database running after the extraction:

```
dbxtract -xx -ac "SERVER=remotedb;UID=DBA;PWD=passwd" -c  
"SERVER=consoldb;UID=DBA;PWD=passwd" RemUser
```

4. Connect to the remote database using the following command:

```
dbisql -c "SERVER=remotedb;UID=DBA;PWD=passwd"
```

5. The consolidated database uses the FILE message system, so when `dbxtract` runs, it creates SQL Remote definitions assuming that the remote database also uses the FILE message system. Since the remote database will use the HTTP message system, execute the following statements on the remote database to reset the FILE message system for this remote database:

```
CREATE REMOTE MESSAGE TYPE FILE ADDRESS '';  
SET REMOTE FILE OPTION public.directory='';  
SET REMOTE FILE OPTION public.debug='';
```

- Execute the following statements on the remote database to configure the HTTP message system for this remote database:

```
CREATE REMOTE MESSAGE TYPE HTTP ADDRESS 'remmsgs';
GRANT CONSOLIDATE TO ConUser TYPE HTTP ADDRESS 'conmsgs';
SET REMOTE HTTP OPTION public.user_name='RemUser';
SET REMOTE HTTP OPTION public.password='passwd';
SET REMOTE HTTP OPTION public.debug='yes';
SET REMOTE HTTP OPTION public.https='no';
SET REMOTE HTTP OPTION public.url='localhost:8033';
COMMIT;
```

- Verify that the employees table in the remote database contains the two rows of data that existed in the consolidated database after the extraction. Execute the following statement to view the contents of the employees table:

```
SELECT * FROM employees;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the values you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	2018-07-06 10:30:02.369
2	Melisa	Boysen	2018-07-06 10:30:02.371

## Results

The remote database is extracted and the FILE message system at the remote database is replaced with the HTTP message system.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server \[page 190\]](#)

**Previous task:** [Lesson 2: Creating the Message Server \[page 193\]](#)

**Next task:** [Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases \[page 198\]](#)

## 1.7.4 Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases

Add data to the consolidated and remote database, run SQL Remote to replicate the changes, and then confirm that the data is consistent in both databases.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. If you are not currently connected to the consolidated database, run the following command:

```
dbisql -c "SERVER=consolddb;UID=DBA;PWD=passwd"
```

2. Execute the following statements on the consolidated database to add additional sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Javier', 'Sporr');  
COMMIT;
```

3. 4. If you are not currently connected to the remote database, then run the following command:

```
dbisql -c "SERVER=remotedb;UID=DBA;PWD=passwd"
```

4. Execute the following statements on the remote database to add additional sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Nelson', 'Kreitzer');  
COMMIT;
```

5. For the consolidated database, run the Message Agent:

```
dbremote -c "SERVER=consolddb;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial  
\conlog1.txt
```

This scans the transaction log of the consolidated database and generates a message for the remote database using the FILE message system. Since the debug message system parameter has been set for the FILE message system in the consolidated database, you can look at the `c:\tutorial\conlog1.txt` file and verify that you see debug messages indicating that messages are being written to the `c:\tutorial\messages\remmsgs` directory. For example:

```
I. 2018-07-06 15:19:27. Processing transactions from active transaction log  
I. 2018-07-06 15:19:27. Sending message to "RemUser"  
(0-0000000000-0000720309-0)  
I. 2018-07-06 15:19:27. sopen "c:\tutorial\messages\remmsgs\conmsgs.0"  
I. 2018-07-06 15:19:27. write "c:\tutorial\messages\remmsgs\conmsgs.0"  
I. 2018-07-06 15:19:27. close "c:\tutorial\messages\remmsgs\conmsgs.0"
```

6. For the remote database, run the Message Agent:

```
dbremote -c "SERVER=remotedb;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial\remlog.txt
```

Using the HTTP message system, this command receives and applies the message that was just generated by the consolidated database. It then scans the transaction log and sends a message back to the consolidated database with the new row that was added in the remote database. Since the debug message system parameter has been set for the HTTP message system in the remote database, you can look at the `c:\tutorial\remlog.txt` file and verify that you see debug messages indicating that the HTTP message system is being used. For example:

```
I. 2018-07-06 15:24:36. Sending message to "ConUser"
(0-0000000000-0000727798-0)
I. 2018-07-06 15:24:36. HTTPWriteMessage "remmsgs.0"
I. 2018-07-06 15:24:36. HTTPWriteMessage: success -- filename "remmsgs.0"
I. 2018-07-06 15:24:37. HTTPDisconnect
```

7. For the consolidated database, run the Message Agent specifying a different log file:

```
dbremote -c "SERVER=consolddb;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial\conlog2.txt
```

This command receives and applies the message that was just generated by the remote database using the FILE-based message system.

8. To verify that the consolidated database contains all four rows of data, execute the following statement to view the contents of the employees table:

```
SELECT * FROM employees
ORDER BY employee_id;
```

The query returns the following data from the employees table, although the `hire_date` column contains the time you inserted the row, and not the values you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	2018-07-06 10:30:02.369
2	Melisa	Boysen	2018-07-06 10:30:02.371
3	Javier	Spoor	2018-07-06 13:52:03.502
102000001	Nelson	Kreitzer	2018-07-06 13:52:14.063

9. Disconnect Interactive SQL from the consolidated database.
10. Verify that the remote database contains all four rows of data by executing the following statement to view the contents of the employees table:

```
SELECT * FROM employees
ORDER BY employee_id;
```

The query returns the following data from the employees table, although the `hire_date` column contains the time you inserted the row, and not the values you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	2018-07-06 10:30:02.369

employee_id	first_name	last_name	hire_date
2	Melisa	Boysen	2018-07-06 10:30:02.371
3	Javier	Spoor	2018-07-06 13:52:03.502
102000001	Nelson	Kreitzer	2018-07-06 13:52:14.063

11. Disconnect Interactive SQL from the remote database.

## Results

Data is added to the consolidated and remote database, the changes are replicated, and the data consistency has been verified.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server \[page 190\]](#)

**Previous task:** [Lesson 3: Creating the Remote Database \[page 196\]](#)

**Next task:** [Lesson 5: Finishing up \[page 200\]](#)

## 1.7.5 Lesson 5: Finishing up

Shut down the three database servers you started in this tutorial.

## Prerequisites

You must have the SERVER OPERATOR system privilege to shut down a network server. No privilege is required to shut down a personal server.



## Procedure

1. Run the following command to shut down the remote database:

```
dbstop -y -c "SERVER=remotedb;UID=DBA;PWD=passwd"
```

2. Run the following command to shut down the message server database:

```
dbstop -y -c "SERVER=msgsrvdb;UID=DBA;PWD=passwd"
```

3. Run the following command to shut down the consolidated database:

```
dbstop -y -c "SERVER=consolddb;UID=DBA;PWD=passwd"
```

## Results

The three database servers are shut down.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server \[page 190\]](#)

**Previous task:** [Lesson 4: Adding and Replicating Data in the Consolidated and Remote Databases \[page 198\]](#)

## 1.8 Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server

In this tutorial you learn how to set up a SQL Remote replication system that uses the Relay Server to forward HTTP traffic to the consolidated database.

### Prerequisites

This tutorial refers to a user, DBA, that must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

The consolidated database uses the FILE message system to replicate changes, while the remote database uses the HTTP message system to replicate changes.

In this tutorial, you:

- Create a consolidated SQL Anywhere database and a remote SQL Anywhere database that contains all the data in the consolidated database.
- Create a directory structure to store the messages generated by SQL Remote. The consolidated database accesses the files using the FILE message system, while the remote database uses the HTTP message system.
- Configure an existing Relay Server to forward HTTP traffic to the consolidated database.
- Configure the consolidated database to act as the message server for the HTTP message system and accept forwarded HTTP traffic from the Relay Server.
- Create a remote database that sends messages using the HTTP messaging system.
- Replicate data between the consolidated and remote databases.

1. [Lesson 1: Creating the Consolidated Database \[page 202\]](#)

Create the directories needed to store the databases and their transactions logs, as well as the directory structure for the messages. You also define the schema of the consolidated database, including creation of the remote user and the publication and subscription needed to replicate data.

2. [Lesson 2: Configuring the Relay Server \[page 205\]](#)

Modify the configuration of the Relay Server to forward HTTP requests to the back-end SQL Anywhere database.

3. [Lesson 3: Configuring the Consolidated Database to Act as the Message Server \[page 207\]](#)

Configure the consolidated database to act as the message server for the HTTP message system

4. [Lesson 4: Creating the Remote Database \[page 209\]](#)

Extract the remote database and replace the FILE message system at the remote database with the HTTP message system.

5. [Lesson 5: Adding and Replicating Data in the Consolidated and Remote Databases \[page 211\]](#)

Add data to the consolidated and remote database, run SQL Remote to replicate the changes, and then confirm that the data is consistent in both databases.

6. [Lesson 6: Cleaning Up \[page 213\]](#)

Shut down the Relay Server Outbound Enabler and the consolidated and remote databases.

## 1.8.1 Lesson 1: Creating the Consolidated Database

Create the directories needed to store the databases and their transactions logs, as well as the directory structure for the messages. You also define the schema of the consolidated database, including creation of the remote user and the publication and subscription needed to replicate data.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Context

When SQL Remote runs against the consolidated database, it uses the FILE message system to send and receive messages, but the remote database uses the HTTP message system.

For the purposes of this tutorial, the name of the computer where the consolidated database (and message server), is running is named `machine_cons`.

### i Note

For secure data transfer, use the HTTPS protocol. HTTPS requires certificates, so this tutorial uses HTTP to simplify the steps.

## Procedure

1. Create the following directories to hold the consolidated database and the remote database:
  - `c:\tutorial`
  - `c:\tutorial\cons`
  - `c:\tutorial\rem`
2. Create the following directories to hold the message files generated by the consolidated database and the remote database:
  - `c:\tutorial\messages`
  - `c:\tutorial\messages\cons`
  - `c:\tutorial\messages\rem`
3. From the `c:\tutorial\cons` directory, run the following command to create the consolidated database (`cons`):

```
dbinit -dba DBA,passwd cons.db
```

4. Start the consolidated database:

```
dbsrv17 -n cons c:\tutorial\cons\cons.db -xs http(port=8033)
```

`-xs http(8033)` is required on the command line, because this is the database server that will be accepting HTTP requests from the remote database and accessing the messages files that exist in the `c:\tutorial\messages` directory. No web services are defined at this time. In this lesson, you only start the personal database server, so only SQL Remote processes on this computer are able to communicate with the message server using HTTP. In a production environment, you would typically use the network server so that SQL Remote processes on other computers would also have access to the web services. You have started a network server in this lesson and named it `cons`. If there is another database server in your network already running with this name, you must choose a different name for the network server and modify the connection strings in the remainder of this tutorial to use the alternative name.

5. Using Interactive SQL, connect to the consolidated database (`cons`) as a user with the `SYS_REPLICATION_ADMIN_ROLE` system role:

```
dbisql -c "UID=DBA;PWD=passwd;SERVER=cons;DBN=cons"
```

- To set the global database ID for the consolidated database (cons), execute the following statement (the global database ID is needed so that distinct primary keys are chosen for all databases when using the GLOBAL AUTOINCREMENT default):

```
SET OPTION public.global_database_id=0;
```

- The schema for the database in this tutorial consists of a single table and all the columns and rows from the table replicate to every remote user. Execute the following statements on the consolidated database (cons) to create the single table in the database:

```
CREATE TABLE employees (
    employee_id BIGINT NOT NULL DEFAULT GLOBAL AUTOINCREMENT(1000000)
    PRIMARY KEY,
    first_name VARCHAR(128) NOT NULL,
    last_name VARCHAR(128) NOT NULL,
    hire_date TIMESTAMP NOT NULL DEFAULT TIMESTAMP
);
```

- Execute the following statements on the consolidated database (cons) to add sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Kelly', 'Meloy');
INSERT INTO employees (first_name, last_name) VALUES ('Melisa', 'Boysen');
COMMIT;
```

- Execute the following statement on the consolidated database (cons) to confirm that the table was created and populated with data:

```
SELECT * FROM employees;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the values you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	4/28/2015 11:17 AM
2	Melisa	Boysen	4/28/2015 11:17 AM

- Execute the following statements to create the user cons that has PUBLISH privileges:

```
CREATE USER cons;
GRANT PUBLISH TO cons;
```

- For performance reasons, the HTTP message system is only used at the remote database, and not at the consolidated. The following statements configure the use of the FILE-based message system at the consolidated database:

```
CREATE REMOTE MESSAGE TYPE FILE ADDRESS 'cons';
SET REMOTE FILE OPTION public.directory='c:\\tutorial\\messages';
SET REMOTE FILE OPTION public.debug='yes';
```

- Execute the following statements to create the remote user rem with a password that has the REMOTE privilege and uses the FILE message system with the specified file path:

```
CREATE USER rem IDENTIFIED BY passwd;
GRANT REMOTE TO rem TYPE FILE ADDRESS 'rem';
```

13. A publication describes the set of data to be replicated. Create a publication named `pub_employees` that replicates all rows of the `employees` table. You subscribe a user to a publication by creating a subscription.

```
CREATE PUBLICATION pub_employees ( TABLE employees );  
CREATE SUBSCRIPTION TO pub_employees FOR rem;
```

14. Disconnect from Interactive SQL.

## Results

The directories needed to store the databases and their transaction logs are created, as well as the directory structure for the messages. The schema of the consolidated database is defined, including creation of the remote user and the publication and subscription needed to replicate data.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server \[page 201\]](#)

**Next task:** [Lesson 2: Configuring the Relay Server \[page 205\]](#)

## Related Information

[-xs Database Server Option](#)

## 1.8.2 Lesson 2: Configuring the Relay Server

Modify the configuration of the Relay Server to forward HTTP requests to the back-end SQL Anywhere database.

## Prerequisites

You must have the `SYS_REPLICATION_ADMIN_ROLE` system role.

This tutorial is based on the following assumptions::

1. The name of the computer where the Relay Server is running is called `machine_iis`, and is a Microsoft Windows 2008 Server R2 running Microsoft IIS 7.5.
2. The version of the Relay Server installed on `machine_iis` is SQL Anywhere version 17.
3. The setup instructions for the Relay Server have been followed exactly as outlined in the documentation.
4. You have deployed the Relay Server components to Microsoft IIS 7.5 on Windows Server 2008/Windows Server 2008 R2.

## Context

Configuring the Relay Server relies only on modifying the `rs.config` file used by the Relay Server.

## Procedure

1. Modify the `rs.config` file on `machine_iis` to add entries for the back-end SQL Anywhere database server that acts as the message server:

```
[backend_farm]
id=srhttp_tutorial_farm
description=SQL Anywhere Web Services farm for tutorial
enable=yes
verbosity=5
[backend_server]
id=srhttp_tutorial_server
description=SQL Anywhere Web Services server for tutorial
farm=srhttp_tutorial_farm
enable=yes
verbosity=5
```

2. From the `%WINDIR%\system32\` directory, run the following command to apply the configuration update:

```
iisreset.exe
```

## Results

The configuration of the Relay Server is modified to forward HTTP requests to the back-end SQL Anywhere database.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server \[page 201\]](#)

**Previous task:** [Lesson 1: Creating the Consolidated Database \[page 202\]](#)

**Next task:** [Lesson 3: Configuring the Consolidated Database to Act as the Message Server \[page 207\]](#)

## Related Information

[Relay Server Deployment](#)

[Deploying the Relay Server Components to Microsoft Windows Server \(Command Line\)](#)

## 1.8.3 Lesson 3: Configuring the Consolidated Database to Act as the Message Server

Configure the consolidated database to act as the message server for the HTTP message system

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Context

It is also possible to configure a separate database and database server to act as the message server.

### Procedure

1. Using Interactive SQL, connect as a user with the SYS\_REPLICATION\_ADMIN\_ROLE system role:

```
dbisql -c "SERVER=cons;DBN=cons;UID=DBA;PWD=passwd"
```

2. When a database is first initialized, there are no web services defined to accept HTTP requests from remote users. There are also no definitions to allow the database server to access the directory where the message files are stored. Create these objects by using the sr\_add\_message\_server stored procedure, which takes an optional parameter to specify who owns all the objects. Execute the following statements

on the consolidated database (cons) to define all the objects needed for the message server, and specify that all the objects are owned by the cons user:

```
CREATE ROLE FOR USER cons;
SET REMOTE http OPTION cons.root_directory='c:\\tutorial\\messages';
CALL sr_add_message_server( 'cons' );
COMMIT;
```

3. Some extra configuration is needed when the Relay Server will be forwarding HTTP requests to a back-end SQL Anywhere server. It is possible to set up a high-availability environment for your back-end SQL Anywhere database servers where some nodes are defined as read-only and some are defined as read-write nodes. In this tutorial, there is only a single database server in the system, so you define the database as a read-write node. Execute the following statements on the consolidated database (cons) to define all the objects needed so that the Relay Server recognizes this database server as a read-write node:

```
CREATE PROCEDURE sp_oe_read_status( )
RESULT (doc LONG VARCHAR)
BEGIN
DECLARE res LONG VARCHAR;
SET res='AVAILABLE=TRUE';
CALL sa_set_http_header('Content-Length', LENGTH(res) );
SELECT res;
END;
GO

CREATE SERVICE oe_read_status
TYPE 'raw'
AUTHORIZATION OFF
SECURE OFF
USER DBA
AS CALL sp_oe_read_status();
GO
```

4. Disconnect from Interactive SQL.
5. The Outbound Enabler acts as a channel between the Relay Server and the back-end SQL Anywhere database. On the machine\_cons computer, start the Relay Server Outbound Enabler (RSOE) by running the following command:

```
rsoe2 -cr "host=machine_iis;port=80;url_suffix=/rs/server/rs.dll"
-cs "host=localhost;port=8033"
-f srhttp_tutorial_farm -id srhttp_tutorial_server -v 5 -o rsoe.log
```

## Results

The consolidated database is configured to act as the message server for the HTTP message system.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server \[page 201\]](#)



**Previous task:** [Lesson 2: Configuring the Relay Server \[page 205\]](#)

**Next task:** [Lesson 4: Creating the Remote Database \[page 209\]](#)

## 1.8.4 Lesson 4: Creating the Remote Database

Extract the remote database and replace the FILE message system at the remote database with the HTTP message system.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. From the `c:\tutorial\rem` directory, run the following command to create the remote database (rem):

```
dbinit -dba DBA,passwd rem.db
```

2. In this lesson, you use `dbxtract` to create the remote database. Run the following command to extract the database for the rem user from the consolidated database, and leave the database server for the remote database running after the extraction:

```
dbxtract -xx -ac "SERVER=rem;DBN=rem;DBF=c:\tutorial\rem  
\rem.db;UID=DBA;PWD=passwd;autostop=no" -c  
"SERVER=cons;DBN=cons;UID=DBA;PWD=passwd" rem
```

3. If you are not currently connected to the remote database (rem) from Interactive SQL, then run the following command:

```
dbisql -c "SERVER=rem;DBN=rem;UID=DBA;PWD=passwd"
```

4. The consolidated database uses the FILE message system, so when `dbxtract` runs, it creates SQL Remote definitions assuming that the rem remote database is also using the FILE message system. Since the remote database will use the HTTP message system, execute the following statements on the remote database (rem) to remove the FILE message system for this remote database:

```
DROP REMOTE MESSAGE TYPE FILE;
```

The DROP REMOTE statement also removes the REMOTE FILE OPTION public settings that we defined in Lesson 1.

5. Execute the following statements on the remote database (rem) to configure the HTTP message system for this remote database:

```
CREATE REMOTE MESSAGE TYPE HTTP ADDRESS 'rem';
```

```
GRANT CONSOLIDATE TO cons TYPE HTTP ADDRESS 'cons';
SET REMOTE HTTP OPTION public.user_name='rem';
SET REMOTE HTTP OPTION public.password='passwd';
SET REMOTE HTTP OPTION public.debug='yes';
SET REMOTE HTTP OPTION public.https='no';
SET REMOTE HTTP OPTION public.url='machine_iis:80/rs17/client/rs.dll/
srhttp_tutorial_farm';
COMMIT;
```

6. Verify that the remote database (rem) contains the two rows of data that existed in the consolidated database after the extraction. Execute the following statement to view the contents of the employees table:

```
SELECT * FROM employees;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the data you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	4/28/2015 11:17 AM
2	Melisa	Boysen	4/28/2015 11:17 AM

7. Disconnect from Interactive SQL.

## Results

The remote database is extracted and the FILE message system at the remote database is replaced with the HTTP message system.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server \[page 201\]](#)

**Previous task:** [Lesson 3: Configuring the Consolidated Database to Act as the Message Server \[page 207\]](#)

**Next task:** [Lesson 5: Adding and Replicating Data in the Consolidated and Remote Databases \[page 211\]](#)

## 1.8.5 Lesson 5: Adding and Replicating Data in the Consolidated and Remote Databases

Add data to the consolidated and remote database, run SQL Remote to replicate the changes, and then confirm that the data is consistent in both databases.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. If you are not currently connected to the consolidated database (cons), run the following command:

```
dbisql -c "SERVER=cons;DBN=cons;UID=DBA;PWD=passwd"
```

2. Execute the following statements on the consolidated database (cons) to add additional sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Javier', 'Sporr');  
COMMIT;
```

3. If you are not currently connected to the remote database (rem), then run the following command:

```
dbisql -c "SERVER=rem;DBN=rem;UID=DBA;PWD=passwd"
```

4. Execute the following statements on the remote database (rem) to add additional sample data to the employees table:

```
INSERT INTO employees (first_name, last_name) VALUES ('Nelson', 'Kreitzer');  
COMMIT;
```

5. At the consolidated database (cons), run the Message Agent:

```
dbremote -c "SERVER=cons;DBN=cons;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial  
\cons1.txt
```

This command scans the transaction log of the consolidated database (cons) and generates a message for the remote database (rem) using the FILE message system. Since the debug message system parameter has been set for the FILE message system in the consolidated database, you can look at the `c:\tutorial\cons1.txt` file and verify that you see debug messages indicating that messages are being written to the `c:\tutorial\messages\rem` directory. For example:

```
I. 2011-04-12 09:33:03. Processing transactions from active transaction log  
I. 2011-04-12 09:33:03. Sending message to "rem" (0-0000000000-0000550994-0)  
I. 2011-04-12 09:33:03. sopen "c:\tutorial\messages\rem\cons.0"  
I. 2011-04-12 09:33:03. write " c:\tutorial\messages\rem\cons.0"  
I. 2011-04-12 09:33:03. close " c:\tutorial\messages\rem\cons.0"
```

- At the remote database (rem), run the Message Agent:

```
dbremote -c "SERVER=rem;DBN=rem;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial\rem.txt
```

Using the HTTP message system, this command receives and applies the message that was just generated by the consolidated database. It then scans the transaction log and sends a message back to the consolidated database with the new row that was added in the remote database. Since the debug message system parameter has been set for the HTTP message system in the remote database, you can look at the c:\tutorial\rem.txt file and verify that you see debug messages indicating that the HTTP message system is being used. For example:

```
I. 2011-04-12 09:34:03. Sending message to "cons" (0-0000000000-0000576448-0)
I. 2011-04-12 09:34:03. HTTPWriteMessage "rem.0"
I. 2011-04-12 09:34:03. HTTPWriteMessage: success -- filename "rem.0"
I. 2011-04-12 09:34:03. HTTPDisconnect
```

- At the consolidated database (cons), run the Message Agent:

```
dbremote -c "SERVER=cons;DBN=cons;UID=DBA;PWD=passwd" -qc -v -o c:\tutorial\cons2.txt
```

This command receives and applies the message that was just generated by the remote database using the FILE-based message system.

- To verify that the consolidated database contains all four rows of data, execute the following statement to view the contents of the employees table:

```
SELECT * FROM employees
ORDER BY employee_id;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the values you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	4/28/2015 11:17 AM
2	Melisa	Boysen	4/28/2015 11:17 AM
3	Javier	Spoor	4/28/2015 11:17 AM
102000001	Nelson	Kreitzer	4/28/2015 11:17 AM

- Verify that the remote database (rem) contains all four rows of data by executing the following statement to view the contents of the employees table:

```
SELECT * FROM employees;
```

The query returns the following data from the employees table, although the hire\_date column contains the time you inserted the row, and not the data you see in the following table:

employee_id	first_name	last_name	hire_date
1	Kelly	Meloy	4/28/2015 11:17 AM
2	Melisa	Boysen	4/28/2015 11:17 AM
3	Javier	Spoor	4/28/2015 11:17 AM

employee_id	first_name	last_name	hire_date
102000001	Nelson	Kreitzer	4/28/2015 11:17 AM

10. Disconnect from Interactive SQL.

## Results

Data is added to the consolidated and remote database, the changes are replicated, and the data consistency has been verified.

## Next Steps

Proceed to the next lesson.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server \[page 201\]](#)

**Previous task:** [Lesson 4: Creating the Remote Database \[page 209\]](#)

**Next task:** [Lesson 6: Cleaning Up \[page 213\]](#)

## 1.8.6 Lesson 6: Cleaning Up

Shut down the Relay Server Outbound Enabler and the consolidated and remote databases.

### Prerequisites

You must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

### Procedure

1. Run the following command to shut down the RSOE:

```
rsoe2 -s -cr "host=machine_iis;port=80;url_suffix=/rs17/rs.dll"  
-cs "host=localhost;port=8033"  
-f srhttp_tutorial_farm -id srhttp_tutorial_server
```

2. Run the following command to shut down the consolidated database:

```
dbstop -y -c "SERVER=cons;DBN=cons;UID=DBA;PWD=passwd"
```

Run the following command to shut down the remote database:

```
dbstop -y -c "SERVER=rem;DBN=rem;UID=DBA;PWD=passwd"
```

## Results

The RSOE and the consolidated and remote databases are shut down.

**Task overview:** [Tutorial: Setting up a Replication System Using the HTTP Message System Via the Relay Server \[page 201\]](#)

**Previous task:** [Lesson 5: Adding and Replicating Data in the Consolidated and Remote Databases \[page 211\]](#)

## 1.9 SQL Remote Reference

SQL Remote provides utilities, options, procedures, and statements to work with the software.

### In this section:

[SQL Remote Utilities and Options Reference \[page 214\]](#)

This chapter describes the utilities, options and procedures provided by SQL Remote.

[SQL Remote System Tables \[page 249\]](#)

SQL Remote system information is held in the SQL Anywhere catalog. A more comprehensible version of this information is held in a set of system views.

[SQL Remote SQL Statements \[page 250\]](#)

Use the following SQL statements to execute SQL Remote commands:

### 1.9.1 SQL Remote Utilities and Options Reference

This chapter describes the utilities, options and procedures provided by SQL Remote.

#### In this section:

[SQL Remote Message Agent Utility \(dbremote\) \[page 215\]](#)

SQL Remote Message Agent utility (dbremote) sends and applies SQL Remote messages, and maintains the message tracking system to ensure message delivery.

[Extraction Utility \(dbxtract\) \[page 225\]](#)

The Extraction utility extracts a remote database from a consolidated SQL Anywhere database.

[SQL Remote Options \[page 237\]](#)

Replication options are database options that provide control over replication behavior.

[SQL Remote Stored Procedures \[page 238\]](#)

Use the following stored procedures to manage an HTTP messaging system.

[SQL Remote System Procedures \[page 241\]](#)

Use this stored procedure to provide names and arguments to the interface for customizing replication at SQL Remote databases.

## 1.9.1.1 SQL Remote Message Agent Utility (dbremote)

SQL Remote Message Agent utility (dbremote) sends and applies SQL Remote messages, and maintains the message tracking system to ensure message delivery.

To run SQL Remote you must have the SYS\_RUN\_REPLICATION\_ROLE system role.

### ≡ Syntax

```
dbremote [ options ] [ transaction-logs-directory ]
```

## Option

## Description

@data

Reads options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used.

To protect information in the configuration file, you can use the File Hiding utility (dbfhide) to encode the contents of the configuration file.

The environment variable can contain any set of options. For example, the first of the following pair of statements sets an environment variable holding a set of options for a SQL Remote process that starts with a cache size of 4 MB, receives messages only, and connects to a database named *field* on a database server named *myserver*. The SET statement should be entered all on one line:

```
SET envvar=-m 4096 -r
-c
"Server=myserver;DBN=field;UID=sa;PWD
=sysadmin"
dbremote @envvar
```

The configuration file contains line breaks, and can contain any set of options. For example, the following command file holds a set of options for a SQL Remote Message Agent that starts with a cache size of 4 MB, sends messages only, and connects to a database named *field* on a database server named *myserver*:

```
-m 4096
-s
-c
"Server=myserver;DBN=field;UID=sa;PWD
=sysadmin"
```

If this configuration file is saved as `c:\config.txt`, it can be used in a command as follows:

```
dbremote @c:\config.txt
```

-a

Processes the received messages (those in the inbox) without applying them to the database. Used together with -v (for verbose output) and -p (so the messages are not purged), this option can help detect problems with incoming messages. Used without -p, this option purges the inbox without applying the messages, which may be useful if a subscription is being restarted.

-b

Runs in batch mode. In this mode, the SQL Remote Message Agent processes incoming messages, scans the transaction log once, processes outgoing messages, and then stops.



Option	Description
<p><code>-c "keyword=value;..."</code></p>	<p>Specifies connection parameters. If this option is not specified, the environment variable SQLCONNECT is used.</p> <p>For example, the following statement runs dbremote on a database file located at <code>c:\mydata.db</code>, connecting with user ID <i>DBA</i> and password <i>passwd</i> (user <i>DBA</i> must have the SYS_RUN_REPLICATION_ROLE system role):</p> <pre data-bbox="810 600 1382 674">dbremote -c "UID=DBA;PWD=passwd;DBF=c:\mydata.db"</pre> <p>The SQL Remote Message Agent must be run by a user with the SYS_RUN_REPLICATION_ROLE system role.</p> <p>The SQL Remote Message Agent supports the full range of SQL Anywhere connection parameters.</p>
<p><code>-dl</code></p>	<p>Displays messages in the SQL Remote Message Agent window or at a command prompt and in the transaction log file if specified.</p>
<p><code>-ek key</code></p>	<p>Specifies that you want to be prompted at a command prompt for the encryption key for strongly encrypted databases. If you have a strongly encrypted database, you must provide the encryption key to use the database or transaction log in any way, including offline transaction logs. For strongly encrypted databases, you must specify either <code>-ek</code> or <code>-ep</code>, but not both. The command fails if you do not specify a key for a strongly encrypted database.</p>
<p><code>-ep</code></p>	<p>Specifies that you want to be prompted for the encryption key. This option causes a window to appear, in which you enter the encryption key. It provides an extra measure of security by never allowing the encryption key to be seen in clear text. For strongly encrypted databases, you must specify either <code>-ek</code> or <code>-ep</code>, but not both. The command fails if you do not specify a key for a strongly encrypted database.</p>
<p><code>-g n</code></p>	<p>Instructs the SQL Remote Message Agent to group transactions containing fewer than <i>n</i> operations together with transactions that follow. The default is twenty operations. Increasing the value of <i>n</i> can speed up processing of incoming messages by doing fewer commits. However, it can also cause deadlock and blocking by increasing the size of transactions.</p>

Option	Description
<code>-l length</code>	<p>Specifies the maximum length of each message to be sent. The <i>size</i> is the amount of memory, in bytes. Use k, m, or g to specify units of kilobytes, megabytes, or gigabytes, respectively. Longer transactions are split into more than one message. The default is 50000 bytes and the minimum length is 10000 bytes.</p> <p><b>⚠ Caution</b></p> <p>The maximum message length must be the same at all sites in an installation.</p> <p>For platforms with restricted memory allocation, the value must be less than the maximum memory allocation of the operating system.</p>
<code>-m size</code>	<p>Specifies a maximum amount of memory to be used by the SQL Remote Message Agent for building messages and caching incoming messages. The <i>size</i> is the amount of memory, in bytes. Use k, m, or g to specify units of kilobytes, megabytes, or gigabytes, respectively. The default is 2048 kilobytes (2 MB).</p> <p>When all remote databases are receiving unique subsets of the operations being replicated, a separate message for each remote database is built up concurrently. Only one message is built for a group of remote users that are receiving the same operations. When the memory being used exceeds the -m value, messages are sent before reaching their maximum size (as specified by the -l option).</p> <p>When messages arrive, they are stored in memory by the SQL Remote Message Agent until they are applied. This caching of messages prevents rereading messages that are out of order from the message system, which may lower performance on large installations. When the memory usage specified using the -m option is exceeded, messages are deleted in a least-recently used fashion.</p>

Option	Description
<code>-ml directory</code>	<p>Specifies the location of offline transaction log mirror files. This option makes it possible for dbremote to delete old transaction log mirror files when either of the following two circumstances occur:</p> <ul style="list-style-type: none"> <li>the offline transaction log mirror is located in a different directory from the transaction log mirror</li> <li>dbremote is run on a different computer from the remote database server</li> </ul> <p>In a typical setup, the active transaction log mirror and renamed transaction log mirrors are located in the same directory, and dbremote is run on the same computer as the remote database, so this option is not required and old transaction log mirror files are automatically deleted. Transaction logs in this directory are only affected if the <code>delete_old_logs</code> database option is set to a value other than Off.</p>
<code>-O filename</code>	<p>Prints messages to the specified log file. The default is to print output to the screen.</p>
<code>-OS size</code>	<p>Specifies the maximum file size for logging output messages. The <i>size</i> is the amount of memory, in bytes. Use <i>k</i>, <i>m</i>, or <i>g</i> to specify units of kilobytes, megabytes, or gigabytes, respectively. By default, there is no limit, and the minimum limit is 10000 bytes.</p> <p>Before SQL Remote logs output messages to an output log file, it checks the current file size. If the log message will make the file size exceed the specified size, SQL Remote renames the output file to <code>yymmddxx.db</code>, where <i>xx</i> is a number that starts at 00 and continues incrementing (<i>xx</i> can be more than 2 digits long), and <i>yymmdd</i> represents the current year, month, and date.</p> <p>If the SQL Remote Message Agent is running in continuous mode for a long time, this option allows you to manually delete old output log files and free up disk space.</p>
<code>-ot file</code>	<p>Truncates the output log file and then appends output messages to it. The default is to send output to the screen.</p>
<code>-p</code>	<p>Does not purge messages.</p>
<code>-q</code>	<p>Starts the SQL Remote Message Agent with a minimized window. This option applies to Microsoft Windows operating systems only.</p>
<code>-qc</code>	<p>Closes SQL Remote window on completion.</p>

Option	Description
<code>-r</code>	<p>Receives messages. If neither <code>-r</code> or <code>-s</code> is specified, the SQL Remote Message Agent executes both phases. Otherwise, only the indicated phases are executed.</p> <p>The SQL Remote Message Agent runs in continuous mode if it is started with the <code>-r</code> option. To have the SQL Remote Message Agent shut down after receiving messages, use the <code>-b</code> option in addition to <code>-r</code>.</p>
<code>-rd minutes</code>	<p>Specifies the polling frequency for incoming messages. By default, the SQL Remote Message Agent polls for incoming messages every minute. This option (<code>rd</code> stands for <i>receive delay</i>) allows the polling frequency to be configured, which is useful when polling is expensive. The default value is one minute.</p> <p>You can use a suffix of <code>s</code> after the number to indicate seconds, which may be useful if you want frequent polling. For example, the following command polls every thirty seconds.</p>
<code>-ro filename</code>	<pre data-bbox="826 1023 1054 1055">dbremote -rd 30s</pre> <p>The <code>-rd</code> option is often used with the <code>-rp</code> option that sets the number of polls for which the SQL Remote Message Agent waits before requesting that a missing message be re-sent.</p>
<code>-ro filename</code>	<p>Logs remote output to file. This option is for use at consolidated sites. When remote databases are configured to send output log information to the consolidated database, this option writes the information to a file. The option is provided to help administrators troubleshoot errors at remote sites.</p>

Option	Description
<code>-rp number</code>	<p>Specifies the number of receive polls before a message is assumed lost. When running in continuous mode, the SQL Remote Message Agent polls at certain intervals for messages. After polling a set number of times (by default, one), if a message is missing, the SQL Remote Message Agent assumes it has been lost and requests that the message be re-sent. On slow message systems, this behavior can result in many unnecessary resend requests. You can set the number of polls before a resend request is issued using this option to minimize the number of resend requests.</p> <p>The <code>-rp</code> option is often used in conjunction with the <code>-rd</code> option that sets the polling frequency for incoming messages.</p>
<code>-rt filename</code>	<p>Truncates the output log file on startup, and then appends the log output from the remote database to the file. This option is for use at consolidated sites. It is identical to the <code>-ro</code> option except that the file is truncated on startup.</p>
<code>-ru time</code>	<p>Specifies the waiting period to re-scan the log on receipt of a resend.</p> <p>This option controls the <i>resend urgency</i>. This is the time between detection of a resend request and when the SQL Remote Message Agent starts fulfilling the request. Use this option to help the SQL Remote Message Agent collect resend requests from multiple users before rescanning the log. The time unit can be s (seconds), m (minutes), h (hours), or d (days).</p>
<code>-s</code>	<p>Sends messages. If neither <code>-r</code> or <code>-s</code> is specified, the SQL Remote Message Agent executes both phases. Otherwise, only the indicated phases are executed.</p>
<code>-sd time</code>	<p>Controls the delay between polls of the database transaction log. The <code>-sd</code> option is only used when running in continuous mode. The default value is one minute.</p> <p>Controls the <i>send delay</i> that is the time to wait between polls for more transaction log data to send.</p>

Option	Description
-t	<p>Replicates all triggers. If you use this option, you must ensure that the trigger actions are not carried out twice at remote databases, once by the trigger being fired at the remote site, and once by the explicit application of the replicated actions from the consolidated database.</p> <p>To ensure that trigger actions are not carried out twice, you can wrap an IF CURRENT REMOTE USER IS NULL...END IF statement around the body of the triggers.</p>
-ts session-name(session-option=[option-value;...])	<p>Sets up a SQL Remote tracing session. The session name must be <i>logging</i>.</p> <p>All information specified after the <i>-ts logging</i> portion of the option must be specified without any spaces.</p>
-u	<p>Processes only transactions that exist in offline transaction logs. This option prevents the SQL Remote Message Agent from processing transactions since the latest backup. Using this option, outgoing transactions and confirmation of incoming transactions are not sent until they exist in offline transaction logs.</p> <p>Only transactions from renamed logs are processed.</p>
-ud	<p>Runs the SQL Remote Message Agent as a daemon on UNIX and Linux platforms. If you run the SQL Remote Message Agent as a daemon, you must also supply the -o or -ot option to log output information.</p> <p>If you run the SQL Remote Message Agent as a daemon and are using FTP or SMTP message links, you must store the message link parameters in the database because the SQL Remote Message Agent does not prompt the user for these options when running as a daemon.</p> <p>When you start the SQL Remote Message Agent as a daemon, its permissions are controlled by the current user's umask setting. Set the umask value before starting the SQL Remote Message Agent to ensure that it has the appropriate permissions.</p>
-ui	<p>For Linux with X window server support, starts the SQL Remote Message Agent in shell mode if a usable display is not available.</p>

Option	Description
-ux	<p>Opens the SQL Remote Message Agent window on Solaris and Linux.</p> <p>When -ux is specified, dbremote 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, dbremote fails to start. On Microsoft Windows, the SQL Remote messages window appears automatically.</p>
-v	<p>Displays verbose output. This option displays the SQL statements contained in the messages to the messages window and, if the -o or -ot option is used, to a message log file.</p>
-W n	<p>Specifies the number of database worker threads to apply incoming messages. This option is not supported on Microsoft Windows Mobile.</p> <p>The default is zero, which means all messages are applied by the main (and only) thread. A value of 1 (one) would have one thread receiving messages from the message system and one thread applying messages to the database. The maximum number of database worker threads is 50.</p> <p>The -w option makes it possible to increase the throughput of incoming messages with hardware upgrades. Putting the consolidated database on a device that can perform many concurrent operations (a RAID array with a striped logical drive), can improve throughput of incoming messages. Multiple processors in the computer running the SQL Remote Message Agent could also improve throughput of incoming messages.</p> <p>The -w option does not improve performance significantly on hardware that cannot perform many concurrent operations.</p> <p>Incoming messages from a single remote database are never applied on multiple threads. Messages from a single remote database are always applied serially in the correct order.</p>

Option	Description
<code>-x [ size ]</code>	<p>Renames and restarts the transaction log after it has been scanned for outgoing messages. In some circumstances, replicating data to a consolidated database can take the place of backing up remote databases, or renaming the transaction log when the database server is shut down.</p> <p>If the optional <code>size</code> qualifier is supplied, the transaction log is renamed only if it is larger than the specified size. The <code>size</code> is the amount of memory, in bytes. Use <code>k</code>, <code>m</code>, or <code>g</code> to specify units of kilobytes, megabytes, or gigabytes, respectively. The default is 0.</p>
<code>transaction-logs-directory</code>	<p>The directory that contains the transaction log for the SQL Anywhere remote or consolidated database. There is an active transaction log and zero or more transaction log archive files, all of which may be required by <code>dbremote</code> to determine what to replicate. Specify this parameter if the transaction log archive files are in a different directory than the active transaction log.</p> <p>There is a slight performance penalty when using this feature because the database server must do extra work to retrieve the transaction logs.</p>

## Remarks

You can run the SQL Remote Message Agent from your own application by calling into the DBTools library. For more information, see the file `dbrmt.h` in the `%SQLANY17%\SDK\Include\` directory.

The user ID in the SQL Remote Message Agent command must have the `SYS_RUN_REPLICATION_ROLE` system role.

The SQL Remote Message Agent uses several database connections.

### Message system control parameters

SQL Remote uses several registry settings to control aspects of message link behavior.

On Microsoft Windows, the message link control parameters are stored in the registry, at the following location:

```
HKEY_CURRENT_USER\Software\SAP\SQL Remote
```

## Related Information

[SQL Remote Message Agent \(dbremote\) \[page 95\]](#)



[SQL Remote Message Systems \[page 120\]](#)  
[Configuration Files](#)  
[Alphabetical List of Connection Parameters](#)  
[The CURRENT REMOTE USER Special Value \[page 53\]](#)  
[Collecting Errors from the Remote Database \(SQL\) \[page 151\]](#)  
[GRANT ROLE Statement](#)  
[File Hiding Utility \(dbfhide\)](#)  
[GRANT ROLE Statement](#)  
[delete\\_old\\_logs Option \[MobiLink\]\[SQL Remote\]](#)  
[Troubleshooting Performance When Receiving Messages \[page 104\]](#)  
[Troubleshooting Performance When Receiving Messages \[page 104\]](#)  
[-ts mlsrv17 Option](#)  
[Event Trace Data \(ETD\) File Management Utility \(dbmanageetd\)](#)  
[Remote Message Type Control Parameters \[page 124\]](#)

## 1.9.1.2 Extraction Utility (dbxtract)

The Extraction utility extracts a remote database from a consolidated SQL Anywhere database.

Users must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

Syntax	
<code>dbxtract [ options ] [ directory ] subscriber</code>	
Option	Description
<code>@data</code>	<p>Reads options from the specified environment variable or configuration file. If both exist with the same name, the environment variable is used.</p> <p>To protect information in the configuration file, you can use the File Hiding utility (dbfhide) to encode the contents of the configuration file.</p>

Option	Description
<p><code>-ac "keyword=value;..."</code></p>	<p>Connects to the database specified in the connection string to do the reload.</p> <p>You can combine the operation of unloading a database and reloading the results into an existing database using this option.</p> <p>For example, the following command (entered all on one line) loads a copy of the data for the <code>field_user</code> subscriber into an existing database file named <code>field.db</code>.</p> <pre data-bbox="810 680 1385 808"> dbxtract -c "UID=DBA;PWD=passwd;DBF=cons.db" -ac "UID=DBA;PWD=passwd;DBF=field.db" field_user </pre> <p>When you specify this option, no copy of the data is created on disk, so you do not specify an unload directory in the command. This provides greater security for your data, but at some cost for performance.</p> <p>When you specify this option, password values are unloaded and reloaded into the resulting database.</p>
<p><code>-ae</code></p>	<p>Use this option with the <code>-an</code> or <code>-ac</code> option to specify that the extraction operation proceed to completion even if errors are encountered. When <code>-ae</code> is specified, <code>dbxtract</code> attempts to execute all SQL statements. If any statements result in an error, then they are logged in the <code>unprocessed.sql</code> script file. After attempting to generate a remote database that encountered errors, review and fix the errors and rebuild the remote database from scratch (recommended), or fix them in the <code>unprocessed.sql</code> script file and apply them to the remote database if required.</p>
<p><code>-al filename</code></p>	<p>Specifies the transaction log file name for the new database if using the <code>-an</code> option.</p>

Option	Description
<p><code>-an database</code></p>	<p>Creates a database file with the same settings as the database being extracted and automatically reloads it.</p> <p>You can combine the operations of unloading a database, creating a new database, and loading the data using this option.</p> <p>For example, the following command (entered all on one line) creates a new database file named <code>field.db</code> and copies the schema and data for the <code>field_user</code> subscriber of <code>cons.db</code> into it.</p> <pre data-bbox="810 719 1386 824">dbextract -c "UID=DBA;PWD=passwd;DBF=cons.db" -an field.db -ae field_user</pre> <p>When you specify this option, no copy of the data is created on disk, so you need not specify an unload directory in the command.</p> <p>When you specify this option, password values are unloaded and reloaded into the resulting database.</p>
<p><code>-ap size [k]</code></p>	<p>Sets the page size of the new database. This option is ignored unless <code>-an</code> is used. The page size for a database can be (in bytes) 2048 (deprecated), 4096, 8192, 16384, or 32768, with the default being the page size of the original database. Use <i>k</i> to specify units of kilobytes (for example, <code>-ap 4k</code>). If there are already databases running on the database server, the server's page size (set with the <code>-gp</code> option) must be large enough to handle the new page size.</p>
<p><code>-b</code></p>	<p>Does not start subscriptions. If this option is specified, subscriptions at the consolidated database (for the remote database) and at the remote database (for the consolidated database) must be started explicitly using the <code>START SUBSCRIPTION</code> statement for replication to begin.</p>

Option	Description
-c "keyword=value;..."	<p>Specifies database connection parameters, in a string.</p> <p>The user ID should have the SYS_RUN_REPLICATION_ROLE system role to ensure that the user has privileges on all the tables in the database.</p> <p>For example, the following statement (entered all on one line) extracts a database for remote user ID joe_remote from the sample database running on the sample_server database server, connecting as user ID DBA with password passwd. The data is unloaded into the \temp\extract directory.</p> <pre data-bbox="826 763 1366 837">dbextract -c "Server=sample_server;DBN=demo;UID=DBA;PWD=sql" \temp\extract joe_remote</pre> <p>If connection parameters are not specified, connection parameters from the SQLCONNECT environment variable are used, if set.</p>
-d	<p>Extracts data only. If this option is specified, the schema definition is not unloaded and publications and subscriptions are not created at the remote database. This option is used when a remote database already exists with the proper schema, and only needs to be filled with data.</p>
-dt directory	<p>Specifies the temporary directory used during the extraction. The directory must exist and is relative to the client. It specifies the temporary directory for the unprocessed.sql file, if it cannot be written to its file path. If this option is not specified, the location for the temporary directory is specified by one of the SATMP, TMP, TMPDIR, or TEMP environment variables.</p>

## Option

## Description

`-ea alg`

Specifies the encryption algorithm for the new database. This option allows you to choose a strong encryption algorithm to encrypt your new database. You can choose either AES (the default) or AES\_FIPS for the FIPS-certified algorithm. AES\_FIPS uses a separate library and is not compatible with AES.

For greater security, specify AES or AES256 for 128-bit or 256-bit strong encryption, respectively. Specify AES\_FIPS or AES256\_FIPS for 128-bit or 256-bit FIPS-certified encryption, respectively. For strong encryption, you must also specify the `-ek` or `-ep` option.

To create a database that is not encrypted, specify `-ea none`, or do not include the `-ea` option (and do not specify `-e`, `-et`, `-ep`, or `-et`).

If you do not specify the `-ea` option, the default behavior is as follows:

- `-ea NONE`, if `-ek`, `-ep`, or `-et` is not specified
- `-ea AES`, if `-ek` or `-ep` is specified (with or without `-et`)
- `-ea SIMPLE`, if `-et` is used without `-ek` or `-ep`

Algorithm names are case insensitive.

FIPS-certified encryption requires a separate license.

`-ek key`

Specifies the encryption key for the new database. This option allows you to create a strongly encrypted database by specifying an encryption key directly in the command. The algorithm used to encrypt the database is AES or AES\_FIPS as specified by the `-ea` option. If you specify the `-ek` option without specifying `-ea`, the AES algorithm is used.

### Caution

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

Option	Description
-ep	<p>Prompts for the encryption key for the new database. This option specifies that you want to create a strongly encrypted database by typing the encryption key in a window. This provides an extra measure of security by never allowing the encryption key to be seen in clear text.</p> <p>You must input the encryption key twice to confirm that it was entered correctly. If the keys don't match, the initialization fails.</p>
-er	<p>Removes encryption from encrypted tables during an unload procedure.</p> <p>When extracting from a database that has table encryption enabled, you must specify either -er or -et to indicate whether the new database has table encryption enabled, otherwise you get an error when attempting to load the data into the new database.</p> <p>The following command (entered all on one line) extracts a database (<code>cons.db</code>) that has encrypted tables, into a new database (<code>field.db</code>) that does not have table encryption enabled, removing encryption from any encrypted tables.</p> <pre data-bbox="810 1144 1382 1238">dbxtract -an field.db -er -c "UID=DBA;PWD=passwd;DBF=cons.db;DBKEY =29bN8cj1z" field_user</pre>

Option	Description
<p><i>-et</i></p>	<p>Enables database table encryption in the new database (-an or -ar must also be specified). If you specify the -et option without the -ea option, the AES algorithm is used. If you specify the -et option, you must also specify -ep or -ek. You can change the table encryption settings for the new database to be different from those of the database you are unloading.</p> <p>When rebuilding a database that has table encryption enabled, you must specify either -er or -et to indicate whether the new database has table encryption enabled, otherwise you get an error when attempting to load the data into the new database.</p> <p>The following example (entered all on one line) unloads a database (<code>cons.db</code>) into a new database (<code>field.db</code>) that has table encryption enabled, and uses AES_FIPS encryption with the key 34jh.</p> <pre data-bbox="813 952 1380 1086">dbextract -an field.db -et -ea AES_FIPS -ek 34jh -c "UID=DBA;PWD=passwd;DBF=cons.db" field_user</pre>
<p><i>-f</i></p>	<p>Extracts fully qualified publications. Usually you do not need to extract fully qualified publication definitions for the remote database, since it typically replicates all rows back to the consolidated database.</p> <p>However, you may want fully qualified publications for multi-tier setups or for setups where the remote database has rows that are not in the consolidated database.</p>

Option	Description
-g	<p><b>Materialized views</b></p> <p>By default, materialized views defined as MANUAL REFRESH are not initialized after a reload. If you want these materialized views to be initialized as part of the reload process, specify the -g option. Specifying -g causes the database server to execute the sa_refresh_materialized_views system procedure.</p> <p>When deciding whether to use the -g option, consider that initializing all materialized views may cause the reload process to take significantly longer to complete. However, not using the -g option means that the first query that attempts to use an uninitialized materialized view must wait while the database server initializes the view, which may cause an unexpected delay. If you do not use the -g option, you can also manually initialize materialized views after the reload completes.</p> <p><b>Text indexes</b></p> <p>By default, text indexes defined as MANUAL REFRESH are not initialized after a reload. If you want the text indexes initialized as part of the reload process, specify the -g option. Specifying -g causes the database server to execute the sa_refresh_text_indexes system procedure.</p>
-ii	<p>Performs an internal unload and internal reload. Using this option forces the reload script to use the internal UNLOAD and LOAD TABLE statements rather than the Interactive SQL OUTPUT and INPUT statements to unload and load data, respectively. This combination of operations is the default behavior.</p> <p>External operations take the path of the data files relative to the current working directory of dbxtract, while internal statements take the path relative to the database server.</p>
-ix	<p>Performs an internal unload and external reload. Using this option forces the reload script to use the internal UNLOAD statement to unload data, and the Interactive SQL INPUT statement to load the data into the new database.</p> <p>External operations take the path of the data files relative to the current working directory of dbxtract, while internal statements take the path relative to the database server.</p>



Option	Description
<code>-kdi number</code>	Specifies the number of times that the encryption key is hashed. Specify a whole number between 1 and 1000. The default value is 2, which results in the hash being applied 2000 times. The higher the number, the better security. If you are running the database on a slow computer, then a very high number of iterations could result in the database taking longer to start (once the database is running, there is no performance impact). This option must be specified with either the <code>-ek</code> or the <code>ep</code> option, as well as either the <code>-ar</code> or the <code>-an</code> option.
<code>-l level</code>	Performs all extraction operations at specified isolation level. The default setting is an isolation level of 0. If you are extracting a database from an active database server, you should run it at isolation level 3 to ensure that data in the extracted database is consistent with data on the database server. Increasing the isolation level may result in large numbers of locks being used by the Extraction utility ( <code>dbxtract</code> ), and may restrict database use by other users.
<code>-n</code>	Extracts the schema definition only. With this definition, none of the data is unloaded. The reload file contains SQL statements to build the database schema only. You can use the <code>SYNCHRONIZE SUBSCRIPTION</code> statement to load the data over the messaging system. Publications, subscriptions, <code>PUBLISH</code> , and <code>SUBSCRIBE</code> privileges are part of the schema. <pre>dbxtract -c "UID=DBA;PWD=passwd;DBF= \cons\cons.db" -n -r "\remote \reload.sql" "\remote" field_user</pre>
<code>-nl</code>	Extracts the structure (the same behavior as the <code>-n</code> option), but the resulting <code>reload.sql</code> file also includes <code>LOAD TABLE</code> or <code>INPUT</code> statements for each table. No user data is extracted when this option is used. When you specify <code>-nl</code> , you must also include a data directory so that the <code>LOAD/INPUT</code> statements can be generated, even though no files are written to the directory. This option allows you to generate a reload script without unloading data. You can extract the data by specifying <code>-d</code> . If a database contains a table whose data should not be unloaded, you can avoid unloading the data for that table by using <code>dbxtract -d -e table-name</code> .
<code>-o filename</code>	Outputs messages to the specified log file.
<code>-p character</code>	Specifies an escape character. The default escape character ( <code>\</code> ) can be replaced by another character using this option.

Option	Description
<code>-q</code>	Operates quietly: does not display messages or show windows. When this option is specified, <code>-y</code> must also be specified or the operation fails.  This option is available only for the command line utility.
<code>-r reload-file</code>	Modifies the name and directory of the generated reload SQL script file. The default is <code>reload.sql</code> in the current directory. The directory is relative to the current directory of the client application, not the database server.
<code>-ru unprocessed-file</code>	Modifies the name and directory of the generated unprocessed SQL statements file. The default is <code>unprocessed.sql</code> in the current directory. The directory is relative to the current directory of the client application, not the database server.
<code>-u</code>	Does not order data during the unload. By default, the data in each table is ordered by primary key. Unloads are faster with the <code>-u</code> option, but loading the data into the remote database is slower.
<code>-up</code>	Extracts user passwords. By default, unless <code>-an</code> or <code>-ac</code> is specified (extraction with a reload), passwords are not extracted.
<code>-v</code>	Displays verbose messages. The name of the table being unloaded, the number of rows unloaded, and the SELECT statement used.
<code>-xf</code>	Excludes foreign keys. You can use this option if the remote database contains a subset of the consolidated database schema, and some foreign key references are not present in the remote database.
<code>-xh</code>	Excludes procedure hooks.
<code>-xi</code>	Performs an external unload and internal reload. The default behavior for unloading the database is to use the UNLOAD statement, which is executed by the database server. If you choose an external unload, dbxtract uses the OUTPUT statement instead. The OUTPUT statement is executed on the client.  External operations take the path of the data files relative to the current working directory of dbxtract, while internal statements take the path relative to the database server.
<code>-xp</code>	Does not extract stored procedures from the database.
<code>-xt</code>	Does not extract triggers from the database.
<code>-xv</code>	Does not extract views or materialized views from the database.

Option	Description
-XX	<p>Performs an external unload and an external load. Use the OUTPUT statement to unload the data, and the INPUT statement to load the data into the new database.</p> <p>The default unload behavior is to use the UNLOAD statement, and the default loading behavior is to use the LOAD TABLE statement. The internal UNLOAD and LOAD TABLE statements are faster than OUTPUT and INPUT.</p> <p>External operations take the path of the data files relative to the current working directory of dbxtract, while internal statements take the path relative to the database server.</p>
y	Replace the existing SQL script file without confirmation.
directory	Specifies the directory the files are written to. This is not needed if you specify -an or -ac.
subscriber	Specifies the subscriber for whom the database is being extracted.

## Remarks

By default, the Extraction utility (dbxtract) runs at isolation level zero. If you are extracting a database from an active database server, then run it at isolation level 3 to ensure that data in the extracted database is consistent with data on the database server. Running at isolation level 3 may hamper others' turnaround time on the database server because of the large number of locks required. Run the Extraction utility (dbxtract) when the database server is not busy, or run it against a copy of the database.

The Extraction utility (dbxtract) creates a SQL script file and a set of associated data files. The script file can be run against a newly initialized database to create the database objects and load the data for the remote database.

By default, the SQL script file is named `reload.sql`.

If the remote user is a group, then all the user IDs that are members of that group are extracted. This allows multiple users on a remote database with different user IDs, without requiring a custom extraction process.

When using the Extraction utility (dbxtract) or the *Extract Database Wizard* with a version 10.0.0 or later database, the version of dbxtract used must match the version of the database server used to access the database. If an older version of dbxtract is used with a newer database server, or vice versa, an error is reported.

The Extraction utility (dbxtract) and *Extract Database Wizard* do not unload the objects created for the dbo user ID during database creation. Changes made to these objects, such as redefining a system procedure, are lost when the data is unloaded. Any objects created by the dbo user ID since the initialization of the database are unloaded by the Extraction utility (dbxtract), and so these objects are preserved.

## Failed Extractions

If the `-ae` option was specified and a failure occurs during the build of a database using one of the `-ac` or `-an` options, `dbxtract` leaves behind a file named `unprocessed.sql` in the current directory. The `-ru` option changes the path and file name of this file. The file contains all the SQL statements that failed execution. If the file cannot be written to the path, then it is written to a SQL Anywhere temporary directory (refer to `-dt`).

This file gives you the opportunity to correct, remove, or alter the failing statements. Although the `unprocessed.sql` file is created at the start of processing, it is only written if the `-ae` option is used. Using Interactive SQL, you can connect to the new database and run the updated `unprocessed.sql` file. This process allows you to complete the database build without starting over again, which can save considerable time.

When the `unprocessed.sql` file contains statements, `dbxtract` returns a failure error code to allow other tools or scripts to be aware of the failed rebuild.

## Privileges

You must have the `SELECT ANY TABLE` and `ACCESS USER PASSWORD` to specify the `-ac`, `-an`, and `-up` options.

## Example

To automatically extract a remote database:

1. Ensure that the user you are using to connect to the consolidated database has the `SYS_RUN_REPLICATION_ROLE` system role.
2. Run `dbxtract`. Specify the `-ac` option to extract to an existing database or the `-an` option to extract to a new database.

If you specify `-ac`, then create an empty database before running `dbxtract`. For example, the following command creates an empty database named `mydata.db`, and then runs `dbxtract` to populate the newly created database. The newly extracted database is for the remote user named `field_user`.

```
dbinit -dba DBA,passwd \remote\mydata.db
dbxtract -c "UID=DBA;PWD=passwd;DBF=\consolidated\data.db" -ac
"UID=DBA;PWD=passwd;DBF=\remote\mydata.db" -ae field_user
```

Alternatively, use `-an` to perform the same process in a single step:

```
dbxtract -c "UID=DBA;PWD=passwd;DBF=\consolidated\data.db" -an \remote
\mydata.db -ae field_user
```

The new remote database, `mydata.db`, is created with the appropriate schema, remote users, publications, subscriptions, and triggers. By default, the data from the consolidated database is extracted into the remote databases and the subscriptions are started. However, the Extraction utility (`dbxtract`) does not start the SQL Remote Message Agent, so no messages are exchanged.

## Related Information

[Remote Database Extraction \[page 86\]](#)  
[Database Extraction](#)  
[Separately Licensed Components](#)  
[Configuration Files](#)  
[Simple Obfuscation Versus Strong Encryption](#)  
[Initializing a Materialized View](#)  
[@data Database Server Option](#)  
[File Hiding Utility \(dbfhide\)](#)  
[-gp Database Server Option](#)  
[START SUBSCRIPTION Statement \[SQL Remote\]](#)  
[sa\\_refresh\\_materialized\\_views System Procedure](#)  
[sa\\_refresh\\_text\\_indexes System Procedure](#)  
[SYNCHRONIZE SUBSCRIPTION Statement \[SQL Remote\]](#)

### 1.9.1.3 SQL Remote Options

Replication options are database options that provide control over replication behavior.

#### ☞ Syntax

```
SET [ TEMPORARY ] OPTION  
[ userid. | PUBLIC. ]option-name = [ option-value ]
```

Parameter	Description
<code>option-name</code>	The name of the option being changed.
<code>option-value</code>	A string containing the setting for the option.

#### Remarks

These options are used by the SQL Remote Message Agent, and should be set for the user ID specified in the SQL Remote Message Agent command. They can also be set for general public use.

Option	Values	Default
<code>blob_threshold</code> option [SQL Remote]	Integer (in bytes)	256
<code>compression</code> option [SQL Remote]	Integer, from -1 to 9	6
<code>delete_old_logs</code> option [SQL Remote]	On, Off, Delay, <i>n</i> days	Off
<code>external_remote_options</code> option [SQL Remote]	On, Off	Off

Option	Values	Default
qualify_owners option [SQL Remote]	On, Off	On
quote_all_identifiers option [SQL Remote]	On, Off	Off
replication_error option [SQL Remote]	Stored procedure name	(no procedure)
replication_error_piece option [SQL Remote]	Stored procedure name	(no procedure)
save_remote_passwords option [SQL Remote]	On, Off	On
sr_date_format option [SQL Remote]	<i>date-string</i>	yyyy/mm/dd
sr_time_format option [SQL Remote]	<i>time-string</i>	hh:nn:ss.Ssssss
sr_timestamp_format option [SQL Remote]	<i>timestamp-string</i>	yyyy/mm/dd hh:nn:ss.Ssssss
sr_timestamp_with_time_zone_format option [SQL Remote]	<i>timestamp-with-time-zone-string</i>	yyyy/mm/dd hh:nn:ss.Ssssss +hh:nn
subscribe_by_remote option [SQL Remote]	On, Off	On
verify_all_columns option [SQL Remote]	On, Off	Off
verify_threshold option [SQL Remote]	Integer (in bytes)	1000

## 1.9.1.4 SQL Remote Stored Procedures

Use the following stored procedures to manage an HTTP messaging system.

### In this section:

[sr\\_add\\_message\\_server System Procedure \[page 239\]](#)

This procedure defines the web services needed to accept HTTP requests from remote users and also the definitions to allow the database server to access the directory where the message files are stored.

[sr\\_drop\\_message\\_server System Procedure \[page 240\]](#)

This procedure deletes all objects created by sr\_add\_message\_server.

[sr\\_update\\_message\\_server System Procedure \[page 240\]](#)

This procedure needs to be called whenever the SQL Remote definitions in the message server change.

## 1.9.1.4.1 sr\_add\_message\_server System Procedure

This procedure defines the web services needed to accept HTTP requests from remote users and also the definitions to allow the database server to access the directory where the message files are stored.

### ☰ Syntax

```
CALL sr_add_message_server( 'owner' );
```

### Return Value

None. An error is returned if there are issues creating the objects required to define the message server.

### Remarks

When a database is first initialized, none of the web services needed to accept HTTP requests from remote users is defined, and neither are definitions to allow the database server to access the directory where the message files are stored. The creation of these objects is automated with the use of the `sr_add_message_server` stored procedure, which takes an optional parameter to specify who owns all the objects. Object names cannot be duplicated.

### Example

The following statements cause the message server database (msgsrv) to define all the objects needed for the message server, and specify that all the objects are owned by the `cons` user (the consolidated database in this instance).

```
CREATE ROLE FOR USER cons;  
SET REMOTE http OPTION cons.root_directory='c:\\tutorial\\messages';  
CALL sr_add_message_server( 'cons' );  
COMMIT;
```

### Related Information

[sr\\_drop\\_message\\_server System Procedure \[page 240\]](#)

[sr\\_update\\_message\\_server System Procedure \[page 240\]](#)

## 1.9.1.4.2 sr\_drop\_message\_server System Procedure

This procedure deletes all objects created by sr\_add\_message\_server.

≡, Syntax

```
CALL sr_drop_message_server;
```

### Return Value

None. An error is returned if there are issues creating the objects required to define the message server.

### Remarks

This stored procedure is used to delete all objects created by sr\_add\_message\_server.

### Related Information

[sr\\_add\\_message\\_server System Procedure \[page 239\]](#)

[sr\\_update\\_message\\_server System Procedure \[page 240\]](#)

## 1.9.1.4.3 sr\_update\_message\_server System Procedure

This procedure needs to be called whenever the SQL Remote definitions in the message server change.

≡, Syntax

```
CALL sr_update_message_server( 'owner' );
```

### Return Value

None. An error is returned if there are issues creating the objects required to define the message server.



## Remarks

This procedure takes an optional parameter, the user that will own the objects created in the stored procedure.

## Related Information

[sr\\_add\\_message\\_server System Procedure \[page 239\]](#)

[sr\\_drop\\_message\\_server System Procedure \[page 240\]](#)

## 1.9.1.5 SQL Remote System Procedures

Use this stored procedure to provide names and arguments to the interface for customizing replication at SQL Remote databases.

## Notes

Unless otherwise stated, the following conditions apply to event-hook procedures:

- The stored procedures must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.
- The procedure must not commit or rollback operations, or perform any action that performs an implicit commit. The actions of the procedure are automatically committed by the calling application.
- You can troubleshoot the hooks by turning on the SQL Remote Message Agent verbose mode.

### In this section:

[The #hook\\_dict Table \[page 242\]](#)

The SQL Remote Message Agent uses the #hook\_dict table to pass values to hook functions; hook functions use the #hook\_dict table to pass values back to the SQL Remote Message Agent.

[sp\\_hook\\_dbremote\\_begin System Procedure \[page 242\]](#)

Use this system procedure to add custom actions at the beginning of the replication process.

[sp\\_hook\\_dbremote\\_end System Procedure \[page 243\]](#)

Use this system procedure to add custom actions just before the SQL Remote Message Agent exits.

[sp\\_hook\\_dbremote\\_shutdown System Procedure \[page 244\]](#)

Use this system procedure to initiate a SQL Remote Message Agent shutdown.

[sp\\_hook\\_dbremote\\_receive\\_begin System Procedure \[page 245\]](#)

Use this system procedure to perform actions before the start of the receive phase of replication.

[sp\\_hook\\_dbremote\\_receive\\_end System Procedure \[page 245\]](#)

Use this system procedure to perform actions after the end of the receive phase of replication.

[sp\\_hook\\_dbremote\\_send\\_begin System Procedure \[page 246\]](#)

Use this stored procedure to perform actions before the start of the send phase of replication.

[sp\\_hook\\_dbremote\\_send\\_end System Procedure \[page 246\]](#)

Use this stored procedure to perform actions after the end of the send phase of replication.

[sp\\_hook\\_dbremote\\_message\\_sent System Procedure \[page 247\]](#)

Use this stored procedure to perform actions after any message is sent.

[sp\\_hook\\_dbremote\\_message\\_missing System Procedure \[page 247\]](#)

Use this stored procedure to perform actions when the SQL Remote Message Agent has determined that one or more messages is missing from a remote user.

[sp\\_hook\\_dbremote\\_message\\_apply\\_begin System Procedure \[page 248\]](#)

Use this stored procedure to perform actions just before the SQL Remote Message Agent applies a set of messages from a user.

[sp\\_hook\\_dbremote\\_message\\_apply\\_end System Procedure \[page 249\]](#)

Use this stored procedure to perform actions just after the SQL Remote Message Agent has applied a set of messages from a user.

## 1.9.1.5.1 The #hook\_dict Table

The SQL Remote Message Agent uses the #hook\_dict table to pass values to hook functions; hook functions use the #hook\_dict table to pass values back to the SQL Remote Message Agent.

The #hook\_dict table is created immediately before a hook is called using the following CREATE statement:

```
CREATE TABLE #hook_dict(  
NAME VARCHAR(128) NOT NULL UNIQUE,  
value VARCHAR(255) NOT NULL );
```

## 1.9.1.5.2 sp\_hook\_dbremote\_begin System Procedure

Use this system procedure to add custom actions at the beginning of the replication process.

### Rows in #hook\_dict table

Name	Values	Description
<i>send</i>	<i>true or false</i>	Indicates if the process is performing the send phase of replication.
<i>receive</i>	<i>true or false</i>	Indicates if the process is performing the receive phase of replication.

## Remarks

If a procedure of this name exists, it is called when the SQL Remote Message Agent starts.

## Privileges

Hook procedures can be created by any user with the `MANAGE REPLICATION` system privilege. However, to ensure that the hook can access the `#hook_dict` table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the `SELECT ANY TABLE` and `UPDATE ANY TABLE` system privileges.
- Be defined using the `SQL SECURITY INVOKER` clause of the `CREATE PROCEDURE` statement.

### 1.9.1.5.3 `sp_hook_dbremote_end` System Procedure

Use this system procedure to add custom actions just before the SQL Remote Message Agent exits.

#### Rows in `#hook_dict` table

Name	Values	Description
<i>send</i>	<i>true or false</i>	Indicates if the process is performing the send phase of replication.
<i>receive</i>	<i>true or false</i>	Indicates if the process is performing the receive phase of replication.
<i>exit code</i>	<i>integer</i>	A non-zero exit code indicates an error.

## Remarks

If a procedure of this name exists, it is called as the last event before the SQL Remote Message Agent shuts down.

## Privileges

Hook procedures can be created by any user with the `MANAGE REPLICATION` system privilege. However, to ensure that the hook can access the `#hook_dict` table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the `SELECT ANY TABLE` and `UPDATE ANY TABLE` system privileges.
- Be defined using the `SQL SECURITY INVOKER` clause of the `CREATE PROCEDURE` statement.

### 1.9.1.5.4 `sp_hook_dbremote_shutdown` System Procedure

Use this system procedure to initiate a SQL Remote Message Agent shutdown.

#### Rows in `#hook_dict` table

Name	Values	Description
<i>send</i>	<i>true or false</i>	Indicates if the process is performing the send phase of replication.
<i>receive</i>	<i>true or false</i>	Indicates if the process is performing the receive phase of replication.
<i>shutdown</i>	<i>true or false</i>	This row is <i>false</i> when the procedure is called. If the procedure updates the row to <i>true</i> the SQL Remote Message Agent is shut down.

## Remarks

If a procedure of this name exists, it is called when the SQL Remote Message Agent is neither sending nor receiving messages, and permits a hook-initiated shutdown of the SQL Remote Message Agent.

## Privileges

Hook procedures can be created by any user with the `MANAGE REPLICATION` system privilege. However, to ensure that the hook can access the `#hook_dict` table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the `SELECT ANY TABLE` and `UPDATE ANY TABLE` system privileges.
- Be defined using the `SQL SECURITY INVOKER` clause of the `CREATE PROCEDURE` statement.

## 1.9.1.5.5 **sp\_hook\_dbremote\_receive\_begin System Procedure**

Use this system procedure to perform actions before the start of the receive phase of replication.

### **Rows in #hook\_dict**

None

### **Privileges**

Hook procedures can be created by any user with the MANAGE REPLICATION system privilege. However, to ensure that the hook can access the #hook\_dict table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the SELECT ANY TABLE and UPDATE ANY TABLE system privileges.
- Be defined using the SQL SECURITY INVOKER clause of the CREATE PROCEDURE statement.

## 1.9.1.5.6 **sp\_hook\_dbremote\_receive\_end System Procedure**

Use this system procedure to perform actions after the end of the receive phase of replication.

### **Rows in #hook\_dict**

None

### **Privileges**

Hook procedures can be created by any user with the MANAGE REPLICATION system privilege. However, to ensure that the hook can access the #hook\_dict table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the SELECT ANY TABLE and UPDATE ANY TABLE system privileges.
- Be defined using the SQL SECURITY INVOKER clause of the CREATE PROCEDURE statement.

## 1.9.1.5.7 **sp\_hook\_dbremote\_send\_begin System Procedure**

Use this stored procedure to perform actions before the start of the send phase of replication.

### **Rows in #hook\_dict**

None

### **Privileges**

Hook procedures can be created by any user with the MANAGE REPLICATION system privilege. However, to ensure that the hook can access the #hook\_dict table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the SELECT ANY TABLE and UPDATE ANY TABLE system privileges.
- Be defined using the SQL SECURITY INVOKER clause of the CREATE PROCEDURE statement.

## 1.9.1.5.8 **sp\_hook\_dbremote\_send\_end System Procedure**

Use this stored procedure to perform actions after the end of the send phase of replication.

### **Rows in #hook\_dict**

None

### **Privileges**

Hook procedures can be created by any user with the MANAGE REPLICATION system privilege. However, to ensure that the hook can access the #hook\_dict table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the SELECT ANY TABLE and UPDATE ANY TABLE system privileges.
- Be defined using the SQL SECURITY INVOKER clause of the CREATE PROCEDURE statement.

## 1.9.1.5.9 sp\_hook\_dbremote\_message\_sent System Procedure

Use this stored procedure to perform actions after any message is sent.

### Rows in #hook\_dict

Name	Values
<i>remote user</i>	The message destination.

### Privileges

Hook procedures can be created by any user with the MANAGE REPLICATION system privilege. However, to ensure that the hook can access the #hook\_dict table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the SELECT ANY TABLE and UPDATE ANY TABLE system privileges.
- Be defined using the SQL SECURITY INVOKER clause of the CREATE PROCEDURE statement.

## 1.9.1.5.10 sp\_hook\_dbremote\_message\_missing System Procedure

Use this stored procedure to perform actions when the SQL Remote Message Agent has determined that one or more messages is missing from a remote user.

### Rows in #hook\_dict

Name	Values
<i>remote user</i>	The name of the remote user who will have to resend messages.

## Privileges

Hook procedures can be created by any user with the `MANAGE REPLICATION` system privilege. However, to ensure that the hook can access the `#hook_dict` table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the `SELECT ANY TABLE` and `UPDATE ANY TABLE` system privileges.
- Be defined using the `SQL SECURITY INVOKER` clause of the `CREATE PROCEDURE` statement.

### 1.9.1.5.11 `sp_hook_dbremote_message_apply_begin` System Procedure

Use this stored procedure to perform actions just before the SQL Remote Message Agent applies a set of messages from a user.

#### Rows in `#hook_dict`

Name	Values
<i>remote user</i>	The name of the remote user who sent the messages about to be applied.

## Privileges

Hook procedures can be created by any user with the `MANAGE REPLICATION` system privilege. However, to ensure that the hook can access the `#hook_dict` table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the `SELECT ANY TABLE` and `UPDATE ANY TABLE` system privileges.
- Be defined using the `SQL SECURITY INVOKER` clause of the `CREATE PROCEDURE` statement.



## 1.9.1.5.12 sp\_hook\_dbremote\_message\_apply\_end System Procedure

Use this stored procedure to perform actions just after the SQL Remote Message Agent has applied a set of messages from a user.

### Rows in #hook\_dict

Name	Values
<i>remote user</i>	The name of the remote user who sent the messages that have been applied.

### Privileges

Hook procedures can be created by any user with the MANAGE REPLICATION system privilege. However, to ensure that the hook can access the #hook\_dict table, which is used to pass information in and out of hooks, hooks must meet one of the following requirements:

- Be owned by a user with the SELECT ANY TABLE and UPDATE ANY TABLE system privileges.
- Be defined using the SQL SECURITY INVOKER clause of the CREATE PROCEDURE statement.

## 1.9.2 SQL Remote System Tables

SQL Remote system information is held in the SQL Anywhere catalog. A more comprehensible version of this information is held in a set of system views.

You can use the following views to access SQL Remote data:

- SYSARTICLE system view
- SYSARTICLECOL system view
- SYSPUBLICATION system view
- SYSREMOTEOPTION system view
- SYSREMOTEOPTIONTYPE system view
- SYSREMOTETYPE system view
- SYSREMOTEUSER system view
- SYSSUBSCRIPTION system view

## Related Information

[SYSARTICLE System View](#)

[SYSARTICLECOL System View](#)

[SYSPUBLICATION System View](#)

[SYSREMOTEOPTION System View](#)

[SYSREMOTEOPTIONTYPE System View](#)

[SYSREMOTETYPE System View](#)

[SYSREMOTEUSER System View](#)

[SYSSUBSCRIPTION System View](#)

## 1.9.3 SQL Remote SQL Statements

Use the following SQL statements to execute SQL Remote commands:

- ALTER PUBLICATION statement [MobiLink] [SQL Remote]
- ALTER REMOTE MESSAGE TYPE statement [SQL Remote]
- CREATE PUBLICATION statement [MobiLink] [SQL Remote]
- CREATE REMOTE [MESSAGE] TYPE statement [SQL Remote]
- CREATE SUBSCRIPTION statement [SQL Remote]
- CREATE TRIGGER statement
- DROP PUBLICATION statement [MobiLink] [SQL Remote]
- DROP REMOTE MESSAGE TYPE statement [SQL Remote]
- DROP SUBSCRIPTION statement [SQL Remote]
- GRANT CONSOLIDATE statement [SQL Remote]
- GRANT PUBLISH statement [SQL Remote]
- GRANT REMOTE statement [SQL Remote]
- GRANT ROLE statement
- PASSTHROUGH statement [SQL Remote]
- REMOTE RESET statement [SQL Remote]
- REVOKE CONSOLIDATE statement [SQL Remote]
- REVOKE PUBLISH statement [SQL Remote]
- REVOKE REMOTE statement [SQL Remote]
- REVOKE ROLE statement
- SET REMOTE OPTION statement [SQL Remote]
- START SUBSCRIPTION statement [SQL Remote]
- STOP SUBSCRIPTION statement [SQL Remote]
- SYNCHRONIZE SUBSCRIPTION statement [SQL Remote]
- UPDATE statement [SQL Remote]

### In this section:

[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)

Sets a message control parameter for a SQL Remote message system.

## Related Information

[ALTER PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)  
[ALTER REMOTE \[ MESSAGE \] Statement \[SQL Remote\]](#)  
[CREATE PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)  
[CREATE REMOTE \[MESSAGE\] Statement \[SQL Remote\]](#)  
[CREATE SUBSCRIPTION Statement \[SQL Remote\]](#)  
[CREATE TRIGGER Statement](#)  
[DROP PUBLICATION Statement \[MobiLink\] \[SQL Remote\]](#)  
[DROP REMOTE \[MESSAGE\] Statement \[SQL Remote\]](#)  
[DROP SUBSCRIPTION Statement \[SQL Remote\]](#)  
[GRANT CONSOLIDATE Statement \[SQL Remote\]](#)  
[GRANT PUBLISH Statement \[SQL Remote\]](#)  
[GRANT REMOTE Statement \[SQL Remote\]](#)  
[GRANT ROLE Statement](#)  
[PASSTHROUGH Statement \[SQL Remote\]](#)  
[REMOTE RESET Statement \[SQL Remote\]](#)  
[REVOKE CONSOLIDATE Statement \[SQL Remote\]](#)  
[REVOKE PUBLISH Statement \[SQL Remote\]](#)  
[REVOKE REMOTE Statement \[SQL Remote\]](#)  
[REVOKE ROLE Statement](#)  
[SET REMOTE OPTION Statement \[SQL Remote\] \[page 251\]](#)  
[START SUBSCRIPTION Statement \[SQL Remote\]](#)  
[STOP SUBSCRIPTION Statement \[SQL Remote\]](#)  
[SYNCHRONIZE SUBSCRIPTION Statement \[SQL Remote\]](#)  
[UPDATE Statement \[SQL Remote\]](#)

### 1.9.3.1 SET REMOTE OPTION Statement [SQL Remote]

Sets a message control parameter for a SQL Remote message system.

#### ☰ Syntax

```
SET REMOTE message-system OPTION  
[ userid.| PUBLIC.]option-name = option-value
```

```
message-system :  
FILE  
| FTP  
| HTTP  
| SMTP
```

```
option-name :  
common-options  
| file-options  
| ftp-options  
| http-options  
| smtp-options
```

```
common-options :
```

```
debug
| encode_dll
| max_retries
| output_log_send_limit
| output_log_send_now
| output_log_send_on_error
| pause_after_failure
```

```
file-options :
directory
| invalid_extensions
| unlink_delay
```

```
ftp-options :
active_mode
| host
| invalid_extensions
| password
| port
| reconnect_pause
| reconnect_retries
| root_directory
| suppress_dialogs
| user
```

```
http-options :
| certificate
| client_port
| https
| password
| proxy
| reconnect_pause
| reconnect_retries
| root_directory
| url
| user
```

```
smtp-options :
local_host
| pop3_host
| pop3_password
| pop3_port
| pop3_userid
| smtp_authenticate
| smtp_host
| smtp_password
| smtp_port
| smtp_userid
| suppress_dialogs
| top_supported
```

```
option-value : string
```

## Parameters

**userid**

If you do not specify a `userid`, then the current publisher is assumed.

### **common-options**

These options are common to the FILE, FTP, HTTP, and SMTP message systems:

#### **debug**

This parameter is set either to YES or NO. The default is NO. When set to YES, debug output specific to the message system is displayed. This information can be used for troubleshooting problems in the message system.

#### **encode\_dll**

If you have implemented a custom encoding scheme, you must set this to the full path of the custom encoding DLL that you created.

#### **max\_retries**

By default, when SQL Remote is running in continuous mode and an error occurs when accessing the message system, it shuts down after the send and/or received phases. Use this parameter to specify the number of times you want SQL Remote to retry the send and/or receive phases before it shuts down.

#### **output\_log\_send\_limit**

Limits the amount of information that is sent to the consolidated database. The `output_log_send_limit` option specifies the number of bytes at the end of the output log (that is, the most recent entries) that are sent to the consolidated database. The default is 5K.

#### **output\_log\_send\_on\_error**

Sends log information when an error occurs.

#### **output\_log\_send\_now**

When set to YES, sends output log information to the consolidated database. On the next poll, the remote database sends the output log information and then resets the `output_log_send_now` option to NO.

#### **pause\_after\_failure**

This parameter applies when the `max_retries` parameter is specified to a value other than zero and SQL Remote is running in continuous mode. When an error occurs in the message system, this parameter defines the number of seconds SQL Remote waits between retrying the send and/or receive phases.

### **file-options**

These options apply to the FILE message system only:

#### **directory**

The directory under which the messages are stored. This parameter is an alternative to the `SQLREMOTE` environment variable.

#### **invalid\_extensions**

A comma-separated list of file extensions that you do not want the SQL Remote Message Agent (`dbremote`) to use when generating files in the messaging system.

#### **unlink\_delay**

The number of seconds to wait before attempting to delete a file if the previous attempt to delete the file failed. If no value is defined for `unlink_delay`, then the default behavior is set to pause for 1 second

after the first failed attempt, 2 seconds after the second failed attempt, 3 seconds after the third failed attempt, and 4 seconds after the fourth failed attempt.

### **ftp-options**

These options apply to the FTP message system only:

#### **active\_mode**

This parameter controls how SQL Remote establishes the server/client connection. This parameter is set either to YES or NO. The default is NO (passive mode). Passive mode is the preferred transfer mode and the default for the FTP message system. In passive mode, all data transfer connections are initiated by the client, in this case, the message system. In active mode, the FTP server initiates all data connections.

#### **host**

The host name of the computer where the FTP server is running. This parameter can be a host name (such as ftp.sap.com) or an IP address (such as 192.138.151.66).

#### **invalid\_extensions**

A comma-separated list of file extensions that you do not want dbremote to use when generating files in the messaging system.

#### **password**

The password for accessing the FTP host.

#### **port**

The IP port number used for the FTP connection. This parameter is usually not required.

#### **reconnect\_pause**

The time in seconds to pause between each connection attempt. The default setting is 30 seconds. When you set this parameter, only reconnections are affected. The initial connection made by the FTP link is not affected.

#### **reconnect\_retries**

The number of times the link should try to open a socket with the server before failing. The default value is 4. When you set this parameter, only reconnections are affected. The initial connection made by the FTP link is not affected.

#### **root\_directory**

The root directory within the FTP host site that the messages are stored under.

#### **suppress\_dialogs**

This parameter is set to ON or OFF. The default value is OFF. If it is set to ON, the *Connect* window does not appear after failed attempts to connect to the messaging system. Instead, an error is generated.

#### **user**

The user name for accessing the FTP host.

### **http-options**

These options apply to the HTTP message system only:

#### **certificate**

To make a secure (HTTPS) request, a client must have access to the certificate used by the HTTPS server. The necessary information is specified in a string of semicolon-separated keyword=value pairs.

You can use the file keyword to specify the file name of the certificate. You cannot specify a file and certificate keyword together. The following keywords are available:

Keyword	Abbreviation	Description
file		The file name of the certificate
certificate	cert	The certificate itself
company	co	The company specified in the certificate
unit		The company unit specified in the certificate
name		The common name specified in the certificate

Certificates are required only for requests that are either directed to an HTTPS server or can be redirected from a non-secure to a secure server. Only PEM formatted certificates are supported (for example, `certificate='file=filename'`)

To create a certificate name in a SQL Anywhere database:

```
CREATE OR REPLACE CERTIFICATE certificate_name FROM FILE  
'certificate_file';
```

To use the certificate name for an HTTPS message type:

```
SET REMOTE HTTP OPTION user_name.certificate= 'cert_name=certificate_name';
```

### client\_port

Identifies the port number on which SQL Remote communicates using HTTP. It is provided for, and recommended only for, connections through firewalls that filter "outgoing" TCP/IP connections. You can specify a single port number, ranges of port numbers, or a combination of the two. Specifying a low number of client ports could result in SQL Remote being unable to send and receive messages if the operating system has not released the ports in a timely manner after SQL Remote closes the port on a previous run.

### debug

When set to YES, all HTTP commands and responses are displayed in the output log. This information can be used for troubleshooting HTTP support problems. The default is NO.

### https

Specify whether to use HTTPS (`https=yes`) or HTTP (`https=no`).

### password

The message server database password. The password authenticates to third-party HTTP servers and gateways using RFC 2617 Basic authentication.

### proxy\_host

Specifies the URI of a proxy server. For use when SQL Remote must access the network through a proxy server. Indicates that SQL Remote is to connect to the proxy server and send the request to the message server through it.

### reconnect\_pause

The time in seconds to pause between each connection attempt. The default setting is 30 seconds. When you set this parameter, only reconnections are affected. The initial connection made by the FTP link is not affected.

#### **reconnect\_retries**

The number of times the link should try to open a socket with the server before failing. The default value is 4. When you set this parameter, only reconnections are affected. The initial connection made by the FTP link is not affected.

#### **root\_directory**

This HTTP control parameter is ignored when specified at the client side. You define this control parameter in the message server prior to calling the `sr_add_message_server` or `sr_update_message_server` stored procedure. When using the HTTP message system, the address specified for a remote user or publisher can only contain a single subdirectory, and not multiple subdirectories.

#### **url**

Specify the server name or IP address and optionally the port number of the HTTP server being used, separated by a semicolon. If requests are being passed through the Relay Server, you can optionally add a URL extension to indicate which server farm the request should be passed to.

#### **user**

The message server database user ID. Authenticates to third-party HTTP servers and gateways using RFC 2617 Basic authentication.

### **smtp-options**

These options apply to the SMTP message system only:

#### **local\_host**

The name of the local computer. It is useful on computers where SQL Remote is unable to determine the local host name. The local host name is needed to initiate a session with any SMTP server. In most network environments, the local host name can be determined automatically and this entry is not needed.

#### **pop3\_host**

The name of the computer on which the POP host is running. Typically, it is the same name as the SMTP host. It corresponds to the POP3 host field in the SMTP/POP3 login window.

#### **pop3\_password**

The password used to retrieve mail. It corresponds to the password field in the SMTP/POP3 login window.

#### **pop3\_port**

The number of the port on which the POP server is listening. The default is 110. This corresponds to the port field in the SMTP/POP3 login window.

#### **pop3\_userid**

The user ID used to retrieve mail. The POP user ID corresponds to the user ID field in the SMTP/POP3 login window. You must obtain a user ID from your POP host administrator.

#### **smtp\_authenticate**

Determines whether the SMTP link authenticates the user. The default value is YES. Set this parameter to NO to turn off SMTP authentication.



**smtp\_host**

The name of the computer on which the SMTP server is running. It corresponds to the SMTP host field in the SMTP/POP3 login window.

**smtp\_password**

The password for SMTP authentication. By default, this parameter takes the same value as the pop3\_password parameter. The smtp\_password only needs to be set if the user ID is different from that of the POP server.

**smtp\_port**

The number of the port on which the SMTP server is currently listening. The default is 25. This corresponds to the port field in the SMTP/POP3 login window.

**smtp\_userid**

The user ID for SMTP authentication. By default, this parameter takes the same value as the pop3\_userid parameter. The smtp\_userid only needs to be set if the user ID is different from that of the POP server.

**suppress\_dialogs**

This parameter is set to ON or OFF. The default value is OFF. If it is set to ON, the *Connect* window does not appear after failed attempts to connect to the messaging system. Instead, an error is generated.

**top\_supported**

SQL Remote uses a POP3 command called TOP when enumerating incoming messages. The TOP command may not be supported by all POP servers. When you set the top\_supported parameter to NO, SQL Remote uses the RETR command, which is less efficient but works with all POP servers. The default is YES.

## Remarks

The SQL Remote (dbremote) Message Agent saves message system parameters when the user enters them in the message system window when the message system is first used. In this case, it is not necessary to use this statement explicitly. This statement is most useful when preparing a consolidated database for extracting many databases.

The option names are case sensitive. The case sensitivity of option values depends on the option: Boolean values are case insensitive, while the case sensitivity of passwords, directory names, and other strings depend on the case sensitivity of the file system (for directory names), or the database (for user IDs and passwords).

## Privileges

Publishers can set their own options. Otherwise, you must have the SYS\_REPLICATION\_ADMIN\_ROLE system role.

## Side effects

Automatic commit.

## Standards

### ANSI/ISO SQL Standard

Not in the standard.

## Example

The following statement sets the FTP host to `ftp.mycompany.com` for the FTP link for user `Sam_Singer`:

```
SET REMOTE FTP OPTION Sam_Singer.host = 'ftp.mycompany.com';
```

The following statement stops SQL Remote from using the specified file extensions for messages that are generated:

```
SET REMOTE FTP OPTION "Public"."invalid_extensions"='exe,pif,dll,bat,cmd,vbs';
```

The following statement sets the URL to point to the localhost for the HTTP link for user `Sam_Singer`:

```
SET REMOTE HTTP OPTION Sam_Singer.url='localhost:8033';
```

The following statement sets the HTTP URL to point to a Relay Server that forwards the request to the `srhttp` farm:

```
SET REMOTE HTTP OPTION "public"."url"='iis7.company.com:80/rs/client/rs.dll/srhttp';
```

## Related Information

[Custom Encoding Schemes \[page 119\]](#)

[The FTP Message System \[page 127\]](#)

[The FILE Message System \[page 125\]](#)

[The HTTP Message System \[page 130\]](#)

[The SMTP Message System \[page 134\]](#)

[Collecting Errors from the Remote Database \(SQL\) \[page 151\]](#)

[Tutorial: Setting up a Replication System Using the HTTP Message System with a Separate Message Server \[page 190\]](#)

[Remote Message Type Control Parameters \[page 124\]](#)

[SET OPTION Statement](#)



CREATE CERTIFICATE Statement  
SET REMOTE OPTION Statement [SQL Remote]

# Important Disclaimers and Legal Information

## Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
  - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
  - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

## Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

## Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

## Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

## Bias-Free Language

SAP supports a culture of diversity and inclusion. Whenever possible, we use unbiased language in our documentation to refer to people of all cultures, ethnicities, genders, and abilities.



© 2022 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.