APPENDIX A

Enterprise and Anywhere: Differences

About this Appendix

This appendix summarizes the differences between SQL Remote for Adaptive Server Enterprise and for Adaptive Server Anywhere.

This appendix describes the main differences between these versions of the technology.

Contents

Topic	Page
Types of difference	434
Differences in functionality	435
Differences in approach	436
Limitations for Enterprise to Enterprise replication	438

Types of difference

The differences between the versions of the software are of the following kinds:

- **Functionality** Tasks that can be carried out by one of the two versions, but not by the other.
- **Approach** Although a similar result can be obtained, a different approach is required in each version. This includes tasks that are carried out in ways that are superficially different, but which have the same result.
- **Server differences** Tasks associated with SQL Remote, such as backup management, are different for the two servers. These differences are not described here.

This appendix addresses only replication using Adaptive Server Anywhere as remote databases. There are additional limitations if using Adaptive Server Enterprise as remote servers.

Differences in functionality

The major differences in functionality between SQL Remote for Adaptive Server Enterprise (SRE) and SQL Remote for Adaptive Server Anywhere (SRA) are as follows:

- **Long data types** SRE does not support replication of long data types such as **text** or **image** (also called BLOBS).
- **Schema changes** For SRE, schema changes must be made on a **quiet** system. A quiet system means the following:
  - **No transactions being replicated** There can be no transactions being replicated that modify the tables that are to be altered. All transactions that modified tables being altered must be scanned from the transaction log into the stable queue before the schema is altered. This is performed by running the Message Agent normally, or using the `-i -b` switches. After the Message Agent completes, you can make the schema change.
  - **Message Agent shut down** The Message Agent must be shut down when the schema change is being made.
  - **SQL Remote Open Server** If you are using the SQL Remote Open Server, it must be shut down when the schema change is being made.
- **Trigger action replication** In SRE, trigger actions are replicated. In SRA you have the choice of replicating trigger actions, but by default they are not replicated. The replication of trigger actions requires SRE users to ensure that triggers are not fired at remote databases.
- **Platform availability** SRA is available on a wider variety of platforms that SRE, reflecting the platform availability of the two servers.
- **Publication definitions** Publications in SRA can be more selective than those in SRE. For example, in SRA you can use a WHERE clause with any value. In SRE, you can only use IS NULL and IS NOT NULL conditions in the WHERE clause.

Differences in approach

There are some features of SQL Remote that must be approached in a different manner in SRE and SRA.

- **Partitioning tables that do not contain the subscription expression** In SRA, publications can contain subqueries, and these allow tables that do not contain a partition expression to nevertheless be distributed properly among subscribers. In SRE, an additional column must be added to such tables, containing a list of subscribers, and triggers must be written to maintain the column. This column can have a maximum size of 255.

  For descriptions, see the section entitled "Partitioning tables that do not contain the subscription expression", in the chapters "SQL Remote Design for Adaptive Server Anywhere", and "SQL Remote Design for Adaptive Server Enterprise", of the book *Data Replication with SQL Remote*.

- **Conflict resolution** In SRA, conflict resolution is carried out using a special trigger syntax. In SRE, stored procedures must be written to carry out this task.

  For descriptions, see the section entitled "Managing conflicts", in the chapters "SQL Remote Design for Adaptive Server Anywhere", and "SQL Remote Design for Adaptive Server Enterprise", of the book *Data Replication with SQL Remote*.

- **Storing messages before sending** In SRE, a separate table named the **stable queue** is used to hold changes before replication. In SRA, there is no stable queue; instead, the messages are retrieved from current and old transaction log files.

- **Commands** Whereas SQL Remote tasks such as creating publications are carried out using SQL statements in SRA, they are carried out using system stored procedures in SRE.

  The Sybase Central administration tool hides many of these stylistic differences by providing a common look and feel to the administration of each version of SQL Remote.

Adaptive Server Enterprise procedures and Adaptive Server Anywhere statements

In SQL Remote for Adaptive Server Anywhere, SQL statements are used to carry out the tasks that these stored procedures carry out in Adaptive Server Enterprise. The following table lists the SQL Remote procedures, and how they correspond to SQL statements in Adaptive Server Anywhere:

Adaptive Server Enterprise procedure	Corresponding Adaptive Server Anywhere statement
sp_remote_type	CREATE REMOTE MESSAGE TYPE
sp_remote_type	ALTER REMOTE MESSAGE TYPE
sp_drop_remote_type	DROP REMOTE MESSAGE TYPE
sp_grant_remote	GRANT REMOTE
sp_revoke_remote	REVOKE REMOTE
sp_publisher	GRANT PUBLISH
sp_publisher	REVOKE PUBLISH
sp_create_publication	CREATE PUBLICATION
sp_add_article	
sp_add_article_col	
sp_add_article	ALTER PUBLICATION
sp_remove_article	
sp_add_article_col	
sp_remove_article_col	
sp_drop_publication	DROP PUBLICATION
sp_subscription 'create'	CREATE SUBSCRIPTION
sp_subscription 'drop'	DROP SUBSCRIPTION
sp_subscription 'start'	START SUBSCRIPTION
sp_subscription 'stop'	STOP SUBSCRIPTION
sp_subscription 'synchronize'	SYNCHRONIZE SUBSCRIPTION
sp_passthrough_user	PASSTHROUGH FOR USERID
sp_passthrough_subscription	PASSTHROUGH FOR SUBSCRIPTION
sp_passthrough_stop	PASSTHROUGH STOP

Limitations for Enterprise to Enterprise replication

If you wish to use SQL Remote for replication between Adaptive Server Enterprise databases, rather than with Adaptive Server Anywhere remote databases, you should be aware of the following limitations:

- **Database extraction** The extraction utility creates RELOAD.SQL scripts and data files for building Adaptive Server Anywhere remote databases. Setting up remote ASE databases requires an extraction process created by the customer.

  For more information about how to create an extraction process, see "sp_remote procedure" on page 422.

- **Referential integrity errors** Referential integrity is always checked immediately in Adaptive Server Enterprise, while Adaptive Server Anywhere provides the WAIT_FOR_COMMIT option to control when integrity is checked. This presents difficulties when rows move between remote databases, as in territory realignment.

  For example, suppose an **Order** table has a foreign key to a **Customer** table which has a foreign key to a **SalesRep** table. The Customer table is subscribed by sales rep. **The** Order table is also subscribed by sales rep (it has a redundant column maintained by a trigger).

  When a row in **Customer** is updated to point to a new sales rep, a trigger fires to update the sales rep column in **Order**. The update on **Customer** is replicated as a delete to the old rep and an insert to the new rep. Similarly, the triggerred update on **Order** is replicated as a delete to the old rep and an insert to the new rep.

  The problem occurs because SQL Remote replicates the operations in the order they occur, which means the **Customer** row is deleted before the **Order** rows. This causes a referential integrity error.

- **Schema upgrades** Schema upgrades are difficult to manage when both consolidated and remote databases are Adaptive Server Enterprise databases. Passthrough to remote ASE databases is difficult to carry out.

  The problem is due to the need for a quiet system for schema upgrades (see "Differences in functionality" on page 435). Passthrough puts schema upgrade statements into the normal message stream. The operations that precede the schema upgrade (in the same message or a previous message) cannot possibly have been scanned from the transaction log into the stable queue before the schema change takes place.

- **Synchronize subscription** This is not implemented for Adaptive Server Enterprise remote databases.