

SQL Server Replication Process

SQL Server Replication Terminology

An understanding of SQL Server Replication terminology is useful, but not essential, for setting up and maintaining a replicated system. The following is a list of the terms used with SQL Server Replication.

Article

An article can be an entire table, select rows from a table, specified columns from a table or a stored procedure. An article is what is to be replicated. See also publication.

Distribution Agent

The distribution agent is a process that moves transactions and snapshots held in the distribution database to subscribers. See also Snapshot replication and Transaction replication.

Distribution Database

A distribution database tracks the changes (to publications) that need to be replicated (from publishers). A publisher can act as its own distributor, or a remote database server can be used to hold the distribution database.

For transactional replication, the distribution database tracks changes made to publications.

For merge replication it stores only synchronization history.

Distributor

The server that contains the distribution database.

Immediate Transactional Consistency

All sites are guaranteed to always see the same data as the publisher, with no time lag. With immediate transactional consistency each site must simultaneously commit the change. Immediate Transactional Consistency therefore has a performance impact, because of which it is unlikely to be suitable for high performance databases or where replicating over a slow LAN or a WAN.

Merge Agent

The merge agent is a process that merges changes between publication and subscription databases where merge replication is used. See also Merge Replication.

Merge Replication

Type of replication that allows changes to the data to be made at any site. Changes to publications are then merged with the copies of that publication held at the other server(s). Merge replication cannot guarantee transactional consistency because the same or related records can be updated on different servers at the same time. See also snapshot replication and transactional replication.

Publication

The data to be replicated is contained within a publication. A publication can contain a selection of tables and stored procedures. A table included in a publication is called an article. A publication therefore defines the set of data that is to be replicated. Each publication can be replicated as a snapshot publication (using snapshot replication), a transactional publication (using transactional replication) or a merge publication (using merge replication). See also article.

Publisher

A publisher is a server that makes data available for other servers, i.e. it is the server that makes data available for replication.

Pull Subscription

With pull subscription the subscriber asks the publisher for periodic updates. With pull subscription each subscriber can specify the best time for the update. Pull subscription is generally considered the better form of replication when a large number of subscribers are involved, or where the load on the LAN needs to be distributed.

Push Subscription

With push subscription the publisher pushes any changes made to the data out to the subscribers. Typically push subscription is used where changes need to be replicated to the subscribers shortly after the change occurs, i.e. where the replication of changes is required in near real time.