

Sybase Replication Server

Alex Gorelik, Yongdong Wang, Mark Deppe

Sybase Inc., 1650 65th St., Emeryville, CA 94608

Sybase is a leading RDBMS vendor that started by providing OLTP systems for the client-server environment and is currently maturing into an enterprise-wide data management solution provider. As a critical part of this enterprise-wide data management strategy, Sybase Replication Server supports data replication in a distributed environment. In such an environment, the same data may be replicated at multiple sites for quick data access and for high data availability. Replicated data can be maintained using strong consistency or weak consistency algorithms. Algorithms for maintaining strong consistency, such as primary copy and quorum consensus, severely limit data availability during network partition.

Sybase Replication Server implements asynchronous, transactionally consistent log replication using primary copy replication model. Figure 1. shows the basic replication system architecture. Applications update the primary database. The updates are logged and scanned from the primary database log by the Log Transfer Manager (LTM). The LTM passes the updates to the primary Replication Server which determines which replicates are interested in the updates and forwards them to the appropriate replicate Replication Server. The replicate Replication Server applies the updates to the replicate database in the same serial transaction order that was applied at the primary site.

A data replication request is called a subscription. A subscription includes data qualifications similar to the *where*

clause in a SQL *select* command, making fine-granularity data replication possible. A highly scalable subscription rules resolution engine determines which subscriptions are interested in which rows. Some rows can be subscribed to by multiple subscriptions, while other rows can migrate in and out of the subscription data set. Both cases are handled in a transactionally consistent manner.

When a subscription is created, its initial data set needs to be copied to the replicate database. When a subscription is dropped, its data set needs to be deleted from the replicate database. Both operations are performed dynamically and keep the data at the replicates transactionally consistent.

The Replication Server scans the primary database log and continuously propagates the updates using store and forward techniques. At the primary database, the Log Transfer Manager prevents the database log from being truncated before the updates are replicated. Distributed mirroring allows messages to be stored at multiple sites to facilitate disaster recovery. Primary and replicate database dumps are coordinated in order to recover the entire distributed system to a consistent state.

Replicates can update primary data asynchronously using request functions. Requests are submitted at the replicates by executing stored procedures. They are delivered to the primary and applied. The updates caused by these requests at the primary are replicated back to the replicates. Concurrency control is simplified, since it can be performed at the primary.

Sybase Replication Server provides an open system interface to heterogeneous systems. A documented Log Transfer Interface allows foreign applications to submit updates at the primary. At the replicate, a flexible, programmable interface allows mapping of Transact SQL commands to any other language and allows users to assign actions to errors.

The Replication Server is very flexible and powerful. Customers have used it to successfully solve many various enterprise data management problems ranging from providing warm stand-bys to building very sophisticated distributed database applications.

Figure 1. Sybase Replication System Architecture

