

# Oracle's Symmetric Replication Technology and Implications for Application Design

Dean Daniels, Lip Boon Doo, Alan Downing, Curtis Elsbernd, Gary Hallmark, Sandeep Jain, Bob Jenkins, Peter Lim, Gordon Smith, Benny Souder, Jim Stamos

Oracle Server Technology  
and Applications Divisions

Oracle Symmetric Replication extends the asynchronous snapshot replication capabilities of Oracle7 to eliminate primary site or row ownership restrictions imposed by the replication system. Any row of any replicated table can be updated at any time by any database storing the row. Data replication configurations can be based on read-only or updatable snapshots of subsets of master tables, peer to peer replication of complete master tables, or a combination of these two configurations. Symmetric Replication is implemented as a layered architecture that includes data transport, replication program generation, and distributed database administration components. The issues that Symmetric Replication raises for application design are being studied in the context of Oracle Applications products.

Configurations that replicate subsets of tables use snapshots, which are materialized views of remote tables, to select the data to be replicated. Snapshots of single tables without aggregation are refreshed incrementally, and more complex snapshots are refreshed by complete evaluation of the defining query. Symmetric Replication allows simple snapshots to be updatable. In Symmetric Replication, any copy of a table that is not a snapshot is called a master table because it can be used for snapshot refreshes. Modifications to a snapshot are propagated to its master table before the snapshot is refreshed.

Configurations that replicate entire tables can use snapshots, or they can use multiple master tables that exchange updates periodically or on demand. Snapshots can be defined on any master table. Fault tolerance is provided by allowing snapshots to refresh from any master table.

Symmetric Replication uses snapshot refresh and deferred remote procedure calls (RPCs) as its data transport mechanisms. Multiple snapshots can be refreshed as a transactionally consistent group. Deferred RPC is a mechanism for queuing transactions that consist of calls to stored procedures. The queued transactions are executed asynchronously at remote databases. The deferred RPC queues and the differential data used for incremental snapshot refresh are stored in the database and use standard database mechanisms for transactions, recovery, and administration.

In most configurations, the deferred RPC mechanisms are invoked by programs that are generated by the Symmetric Replication program generator. The program generator creates triggers that defer RPCs to propagate table updates to master sites, and procedures to

apply the updates. Because Symmetric Replication allows concurrent updates to different copies of data, the generated replication programs contain logic to detect an update conflict. By default, a conflicting replicated transaction is not applied, and is logged in an error table. The program generators can incorporate application logic for conflict resolution.

Symmetric Replication is administered through a distributed replication catalog interface that consists of new data dictionary views and commands that manipulate the replication configuration and manage replicated database schemas. Replication catalog commands invoke replication program generators, create and refresh snapshots, create new master copies of tables, procedures and other database objects, and manage deferred RPC queues.

Symmetric replication builds on the snapshot technology provided by Oracle7 and many other distributed databases. The implementation uses synchronous distributed database technology including distributed SQL, synchronous RPCs, and distributed transactions. According to the taxonomy defined by Davidson et al., Symmetric Replication uses eventual mutual consistency as a correctness criterion. It provides support for simple pessimistic syntactic approaches that avoid conflicts, such as horizontal partitioning and primary site strategies. Simple optimistic semantic approaches that detect and resolve conflicts (e.g., use the most recent update) are also supported. More complex consistency strategies can be constructed.

The problems of designing applications that use update anywhere resolution and incur few conflicts are being studied using Oracle Applications. Oracle Applications is an integrated suite of business applications including Order Entry, Manufacturing, General Ledger, etc., that contains over 1600 tables or views and 42,500 SQL statements. Preliminary study indicates that Oracle Applications can use Symmetric Replication to share data among functional partitions of the application suite with few conflicts. A few general paradigms can be used to resolve conflicts. Horizontal partitioning of data is also feasible.

Experience with Symmetric Replication has suggested a number of topics for additional research, including incremental view maintenance and work-flow technologies.

Davidson, S.B., Garcia-Molina, H., and Skeen, D. *Consistency in Partitioned Networks*, ACM Computing Surveys, Vol. 17, No 3, Sept 1985.