



Data Replication

Keywords: Redundancy, Backup, Fault tolerance, Accessibility, Performance

Yaacomo Replication Technology

Database replication ensures redundancy and backups and improves fault-tolerance, availability and accessibility. Yaacomo provides the mechanism for a one-way replication that enables data of one Yaacomo server (the primary server) to be stored in or replicated to one or more Yaacomo database servers (one or more secondary servers). Yaacomo replication ensures consistency between the primary and the redundant secondary servers and guarantees ACID properties. Yaacomo realizes replication with the aid of a log-based full audit trail of every transaction. This audit trail can be useful in monitoring database events.

Yaacomo uses a transactional replication model where the secondary servers receive a full initial copy of the database and then receive updates as data changes. All writes and updates must take place on the primary server. Reads, however, may take place on one or more secondary servers. Secondaries are permanently connected to primaries to receive updates.

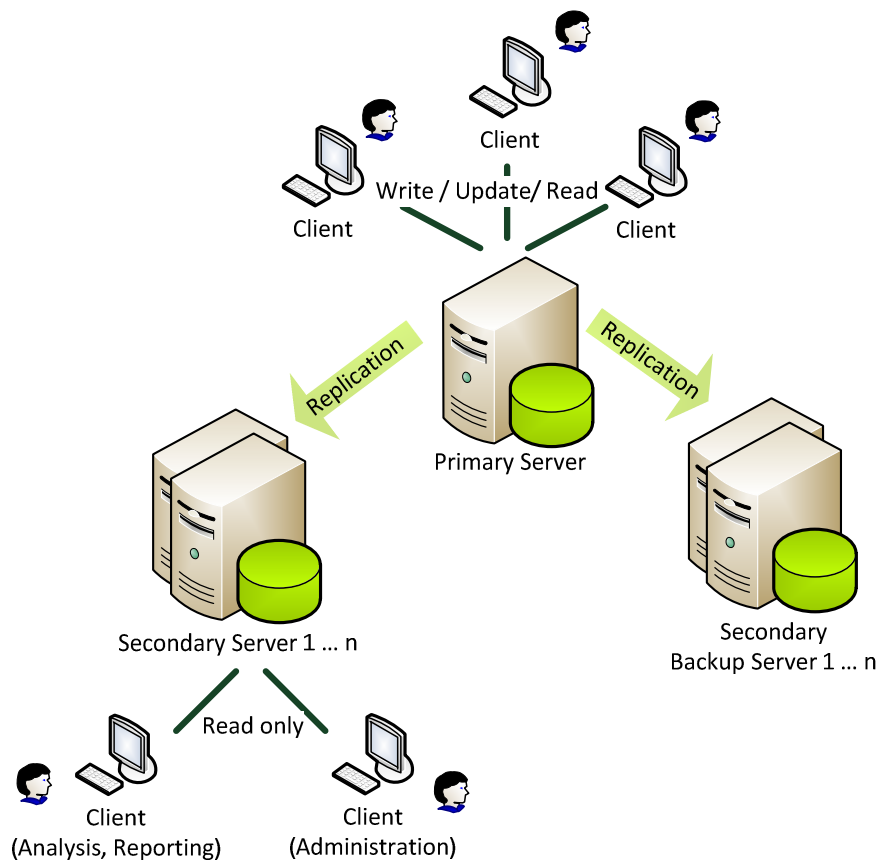
Replication may be either synchronous or asynchronous. This may be configured before starting the database.

Synchronous replication

Synchronous replication writes data to the primary and secondary sites at the same time. A transaction will be successfully ended when and only when the transaction is successfully committed in all secondary servers. Replication ensures that changes, additions, and deletions performed on the data at primary are automatically reflected in the data stored on all other secondary servers. Therefore, every user always sees data that is consistent with the data seen by all the other users.

Asynchronous replication

Asynchronous replication writes data to the primary first and then copies data asynchronously to the secondary servers. With asynchronous replication, there is a delay before the data gets written to the secondary site. Because asynchronous replication is designed to work over longer distances and requires less bandwidth, it is often a better option for disaster recovery. However, asynchronous replication risks a loss of data during a system failure because the data on the secondary is not synchronized with the primary.



Use cases

Fault-Tolerance, Availability

Database replication improves fault-tolerance and provides disaster recovery and high availability. In case of an emergency the secondary database server can be switched to primary to substitute for the original (primary) database.

Scalability

Database replication enables easy scalability of a system by adding new replicas (secondary servers), if the load increases. By spreading the load on multiple database servers the system throughput (response time for reads) can be increased.

Performance

Dedicating a secondary server for large-scale queries can be used to execute analytics queries without affecting the performance of the primary.

Accessibility

Replication is also useful for data distribution. For reading access, you can create a local copy of the data close to your clients that provides fast local access without requiring permanent access to the distant primary.