

Synchronous vs Asynchronous Replication

Native SQL Server AlwaysOn vs
CloudBasic RDS AlwaysOn/Geo-
Replicate for SQL Server Web\SE\EE

CL**UDBASIC**

Native SQL Server AlwaysOn vs CloudBasic RDS AlwaysOn/Geo-Replicate for SQL Server Web\SE\EE

Definition: Synchronous replication requires the bandwidth of a LAN between the servers, possibly with an extended LAN in two geographically remote computer zones. Asynchronous replication can be implemented on a low speed WAN.

The native SQL Server AlwaysOn mirroring replicates synchronously. It can work between two SQL Servers Enterprise in on-premise network or between two EC2 SQL Servers Enterprise in same AWS region. It supports Active-Active fail-over.

CloudBasic RDS AlwaysOn/Geo-Replicate for SQL Server was designed to support primarily hybrid on-premise to AWS RDS or EC2 SQL Server (and vice versa) cross-region replication scenarios. It replicates asynchronously over a WAN connection and tolerates degradation in connectivity (and even short downtimes of the primary or replica SQL Servers) between on-premise and AWS. It supports all versions of SQL Server, including Web Edition, and allows for deploying of a cross-region Active-Passive fail-over.

Please, visit <http://cloudbasic.net/aws/rds/alwayson/dr/> for more information about Active-Passive fail-over.



Synchronous vs. asynchronous replication

The primary difference between synchronous replication and asynchronous replication is the way in which data is written to the replica. Most synchronous replication products write data to primary storage and the replica simultaneously. As such, the primary copy and the replica should always remain synchronized.

In contrast, asynchronous replication products copy the data to the replica after the data is already written to the primary storage. Although the replication process may occur in near-real-time, it is more common for replication to occur on a scheduled basis. For instance, write operations may be transmitted to the replica in batches on a periodic basis (for example, every one minute). In case of a fail-over event, some data loss may occur.

The benefits of asynchronous replication

Asynchronous replication requires substantially less bandwidth than synchronous replication.

It is designed to work over long distances. Since the replication process does not have to occur in real time, asynchronous replication can tolerate some degradation in connectivity.

In contrast, synchronous replication allows failover from primary to secondary data storage to occur nearly instantaneous, to ensure little to no application downtime. However as noted above, it requires the bandwidth of a LAN between the servers, possibly with an extended LAN in two geographically remote computer zones and may also require specialized hardware (depending on the implementation).

Asynchronous replication use cases supported by CloudBasic RDS AlwaysOn/Geo-Replicate



