

Databases High Availability Disaster Recovery (HADR) is a database replication feature that provides a high availability solution for both partial and complete site failures. HADR protects against data loss by replicating data changes from a source database, referred to as the primary, to a target database, and referred to as the standby. A database that does not use HADR is referred to as a standard database.

Without HADR, a partial site failure requires the database management system (DBMS) server to be rebooted. The length of time it takes to restart the database and the machine where it resides is unpredictable. It can take several minutes before the database is brought back to a consistent state and made available. With HADR, the standby database can take over in seconds. Further, one can redirect the clients that were using the original primary database to the standby database by using automatic client reroute or retry logic in the application.

A complete site failure can occur when a disaster, such as a fire causes the entire site to be destroyed. Because HADR uses TCP/IP for communication between the primary and standby databases, they can be situated in different locations. For example, a primary database might be located at the head office in one city, while the standby database is located in a sales office in another city. If a disaster occurs at the primary site, data availability is maintained by having the remote standby database take over as the primary database with full DB2 functionality. After a takeover operation occurs, one can bring the original primary database back up and return it to its primary database status; this is known as failback.

Agitech's High Availability & Disaster Recovery Services include:

- Clustering architecture
- Implementation of Database Clusters (MySQL, Oracle RAC)
- Storage Management Techniques
- Switchover and Failover Techniques
- Load Balancing
- Database Partitioning
- Dynamic Re-Mastering & Extended RAC