



APPLICATION BRIEF

▶▶▶ AVAILABILITY

To be useful, a database needs to be both reliable and available. Durability ensures the persistence of the data in case the database and the server it runs on fail. Availability is the ability to withstand system failures that would normally disrupt the database's ability to function. It not only preserves the data, but availability also plays a crucial role in ensuring that the database process itself is able to survive.

Although availability is often measured in terms of "uptime," it is not a simple percentage or a single decision point. Total availability can never be guaranteed because not all dangers can be predicted in advance.

It is possible to guard against very common issues, such as system failure. It is also possible to prepare for less likely, but also well-understood, dangers, such as large-scale outages and natural disasters. But in the end, availability involves striking a balance between the cost of preventative measures (such as database replication) and the risk to the business associated with possible downtime.

Availability in VoltDB

VoltDB offers several features that let you "dial in" as much or as little availability protection as you like. Specifically, VoltDB provides K-Safety and Network Partition Detection to guard against local hardware failures and Database Replication as protection against larger-scale outages:

- **K-Safety** works by duplicating individual partitions within the database cluster. All copies of the partitions operate simultaneously, ensuring data consistency at all times. If a node fails, the other copies of the affected partitions continue to process requests, keeping the database online until the node can be repaired and rejoin the cluster. You can configure the amount of protection by setting the K-Safety value, which specifies how many nodes can fail without endangering the database as a whole.
- **Network Partition Detection** works in conjunction with K-Safety. In a K-safe cluster, partition detection protects against network outages that would make the cluster think there has been one or more node failures. If the network fails between nodes, the cluster sees that lack of connectivity as a node failure. It is possible that two or more segments of the cluster might attempt to continue separately in this situation. Network Partition Detection ensures that only one segment of the cluster will survive. Partition detection is critical for situations where network fluctuation could impact cluster connectivity.

- **Database Replication** (DR) performs a function similar to K-Safety, but instead of replicating partitions locally, DR replicates the entire cluster remotely. One of the main uses of DR is to provide disaster recovery — the ability to recover from large-scale outages, such as natural disasters, that might take out the entire cluster and its surrounding infrastructure, rather than individual nodes. In such dire situations, it is impossible to keep the primary database running. Instead, DR minimizes downtime by providing a replica, or “hot spare,” that can substitute for the original database in an emergency.

VoltDB’s availability features can be configured separately or in combination to protect against localized outages and wide-scale environmental failures. When combined with durability features, such as command logging, VoltDB provides the maximum protection against the inevitable contingencies faced in today’s business computing environment.

