

DB2 DBMS Availability, Recoverability, and Performance

Jon Capelli, Senior Enterprise Architect 11/16/17

Availability, recoverability, and performance are key factors that must be considered and planned for in any DBMS (Database Management System). Database administrators are tasked with ensuring that these three factors are an integral part of database design, although each one by itself is critical. It doesn't suffice to have databases be recoverable if you can't access them, or if the performance of retrieving critical data is below customer expectations. How many times have you closed your internet browser because the response was too slow and went to another website? It is frustrating for customers when this happens.

Availability refers to having the database and all the tables, indexes, views associated with that database available to the user as much as possible. In today's world of information being available on every device from cell phones to personal computers and everything in between this presents some challenges from a database administration standpoint. Over time, tables within databases get fragmented from heavy workloads of updating, deleting and adding data. This fragmentation can hinder performance. It's necessary to reorganize the data to maintain performance. Reorganization is accomplished by running database utilities against the database. In addition to that workload, there is a need to change the tables to accommodate new business. Over the years DB2 has grown into a DBMS that can support these demands. With every release of DB2, the product becomes more robust to allow online changes and reorganizations to repair fragmentations.

Recoverability of data is another critical piece of database administration. If we lose the ability to recover data, we lose confidence of our customers and eventually their business. There are two different flavors of recoverability: disaster recovery (DR) and local recovery. Disaster recovery refers to a situation where the data center and all the data stored is no longer available due to a catastrophic event. Not just the data, but all the system components need to be backed up offsite in the event of a true disaster. Local recovery refers to data loss due to a hardware failure or application software issues. Years ago, hardware failure was a more common reason for recovery of local data loss. Now, it's usually associated to an application software issue. The good news with the latter is that most of this type of recovery can be accomplished without business interruption. DBAs can access logs to back out data changes that were in error or provide information to assist application areas for the correction of data.

Performance and the relevance of reviewing access paths are the final key component of DBMS design. Now that we discussed the availability and recoverability of data, it's equally important that the data can be accessed and processed in a timely manner to meet the business needs of our customers. SQL (Structured Query Language) and DML (Data Manipulation Language) are used to interface with a relational database to retrieve and manipulate data. In order to satisfy the business needs of today, it's essential that the access to this data is performed like a fine tuned engine.

Availability, recoverability, and performance are all equally important on any DBMS platform; it's hard to say that one is more important than the other. Every new release of DB2 has been making these three key factors more robust and easier to maintain and achieve.