



Db2 z/OS Table and Index Design for High Performance

This seminar will discuss many issues surrounding alternative table and index designs for performance and availability current with 13 for z/OS. We will look at many different alternatives to help design tables and indexes for the best performance possible for the type of application process and how to achieve high availability. There are many issues to consider for high volume insert and update transactions, and a balancing act to be achieved for responsiveness to readers. Table issues such as denormalization, non-column attributes, indicator columns, repeating columns, derived data, aggregated columns, table flipping, identity columns, surrogate keys, horizontal and vertical partitioning. We will also look at issues with designing VLTBs (Very Large Tables) and XVLTBs (eXtremely Large Tables) – the rules have changed! Index issues such as minimizing indexes, using NPSIs or DPSIs, and using indexes for best possible performance will be discussed.

Course Outline:

Alternative Designs

- Design tables and indexes for best performance for application process
- Issues to consider for high volume insert and update transactions
- Performing a balancing act to be achieved for responsiveness to readers
- Designing tables and indexes for maximum availability

Table Design

- Normalization/denormalization
- Non-column attributes
- Indicator columns
- Repeating columns
- Derived data
- Aggregated columns
- Table flipping
- Identity columns/sequence objects
- Surrogate keys
- Horizontal and vertical partitioning
- Compression/on-the-fly compression
- LOBs/inline LOBs
- Volatile tables
- Logging issues/not logging
- Universal tablespaces (PBR and PBG)
- Clone tables/table switching
- Member Cluster/append processing

Temporal tables

Designing VLTBs (Very Large Tables) and XVLTBs (eXtremely Large Tables)

- Issues different from traditional design
- DSSIZE and partition restrictions/options
- Relative and absolute partitioning
- Random keys vs sequential keys
- Designing for high speed inserts
- Designing to minimize maintenance
- Purging and archiving issues
- Archive tables
- Separating active vs inactive data

Index Design

- Challenges with NPSIs
- Challenges with DPSIs
- Uniqueness, clustering and partitioning
- Index page sizes
- Index page splitting/avoidance
- Compression on indexes
- Index on expression
- Hash access
- Include columns
- Disorganized index usage
- Minimizing number of indexes
- Filtering with indexes via SQL