

DB/IQ IA+ Index Administrator

Design the best indexes to maximize overall DB2[®] application performance

Objective

Index

Evaluatior

ADVANTAGES

DB/IQ IA+ recommends the best index candidates based upon your current application SQL and system activity

Impact Analysis via "cost factors" and accumulated CPU consumption indicates index and application efficiency



Improve DB2[®] application performance - let the Optimizer assist you choose the right key columns to index

Examine all SQL associated with a single table, whether Dynamic or Static, capture trouble-some SQL together with real execution statistics

Include non-executed SQL originating from applications that run infrequently for several hours

Clean up the DB2® Catalog by discarding non-required indexes, "same" or similar indexes with overlapping index key columns



Extracts and scans SQL to build a unique Repository DB, providing a quick cross-reference between application, SQL and column usage

Identifies potential index key column candidates automatically

Models Virtual Indexes automatically according to the optimizer and used predicates (option OOX)

Analyzes Index Impact, showing which index, application and SQL profit most and if any deteriorate from the new index

Influences Impact Analysis with QA's run time execution statistics (execution frequency, CPU time etc.)



DB/IQ IA+ Workflow

SOURCE

FACILITIES & FUNCTIONS

BENEFITS

all SQL statements

REPOSITORY & X-REFERENCE DATABASE

Static SQL including: Plans, Packages, MQTs, Views Common Table Expressions, Triggers etc.

- Catalog Index Analysis More than 25 reports including:
 - Non-used indexes
 - Discrepancy conflictions within RUNSTATS values
 - Indexes with overlapping key columns
 - Index FIRSTKEY columns with skewed distribution
 - Foreign keys with incomplete indexes or even without an index
- Query-By-Example

Simply enforce "best" indexes for

- All Dynamic SQL including:
- ERP Applications
- WEB Applications
- Client-Server

- Repository Analysis • Tables & columns found accessed most commonly
- Potential index key column candidates identified

Query-By-Example through Cross-References & Repository

EXECUTION RUNTIME STATISTICS

| For single SQL and packages: • Accumulated CPU time | > | Index Impact | > | Design new or revise existing indexes |
|-------------------------------------------------------------------------------------------|---|----------------------------------------------------------------------------------------------------------------------------------|-----|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Execution frequency User modified values | | Index Usage – Tops & Flops Application Impact via Cost or C consumption | PU | Be guided by the Optimizer "feedback" and predicates used |
| | | Influence by simulating higher execution frequencies | | |
| Interface to "other" system monitors Collect CPU timings and execution frequency | > | Collect Runtime SQL and statistics • From "other" System Monitors • Or Implement QA's Monitor Tra | ace | Improve Application Performance: Revise index design across all applications before production Validate Index Key Columns Evaluate Index Impact Cost- and Access-Path Based Consider execution frequency & CPU time consumed |
| | | Implement QA's RUNSTATS Migration Utilities: • Transfer Production> Test • Generate volume based • Manipulate values | > | Influence Access Path strategy & index usage • RUNSTATS values used by the Optimizer for AP strategy decisions |



