

Chapter 1: Overview

Client Specific Performance Objectives and Examples	1-2
SQL	3
Static SQL	4
Dynamic SQL	5
Deferred Embedded SQL	6
CLI	7
Schemas	8
Tables	9-10
Indexes	11
Keys	12
Constraints	13
Triggers	14
Storage Structures	15
Databases	16
Catalog	17
Views	18
Sequences	19
Routines	20
Locking, Commit, and Rollback	21
Unit of Work	22
Unit of Recovery	23
Packages and Application Plans	24
Coded Character Sets and CCSIDs	25
Java Database Connectivity	26
Indexes	27-28
Privileges	29
Communications Database	30
DB2 Directory	31
Directory Table Spaces	32
Active and Archive Logs	33
BSDS: Bootstrap Data Set	34
Buffer Pools	35
Resource Limit Facility Database	36
Address Spaces	37
Active and Archive Logs	38
BSDS: Bootstrap Data Set	39
Buffer Pools	40
TEMP Database	41

Chapter 2: Database Design: Introduction

Client Specific Performance Objectives and Examples	1-2
Database Design	3
Physical Database Design	4
Steps for Converting the Logical Design	5
Iterative Process	6
Inputs to Design Tuning	7
Design Tuning	8
DB2 Physical Design Goals	9
Never Ending Job	10
Design Goals: Creating Views	11
Determining What Columns to Index	12

CREATE TABLESPACE	23-25
STOGROUP stogroup-name	26
PRIQTY Integer	27
SECQTY Integer	28
ERASE	29
Free Space	30
NUMPARTS Integer	31
LOCKSIZE	32
LOCKMAX	33
COMPRESS	34
SEGSIZE Integer	35
Create Tablespace Example	36

Chapter 6: Creating DB2 Objects: Part 2

Client Specific Performance Objectives and Examples	1-2
DB2 Base Tables and DB2 Temporary Tables	3-5
Implementing Indexes	6
Indexes - Types	7
Unique Indexes	8
Clustering Indexes	9
Partitioning Indexes	10
Secondary Indexes	11
NOT PADDED	12
Indexes for Avoiding Sorts	13
Schemas	14
CREATE TABLE	15-21
Table - Creation	22-23
Table - Creation	23
LIKE Option	24
Storing Tables in the Database	25
Altering a Table	26-27
Altering a Table - Restrictions	28
Dropping a Table	29
NOT NULL, DEFAULT, and GENERATED	30
GENERATED ALWAYS Option	31-32
Constraints	33
Unique Constraints	34-35
Table Check Constraints	36
Views	37-38
Views - Creation	39-40
DROP VIEW	41
Read-only and Updatable Views	42
CHECK OPTION	43
Indexes	44
CREATE INDEX	45
Index - Creation	46-47
ALLOW REVERSE SCANS	48
INDEXES IN	49
Clustering Indexes	50
What Indexes Should be Created	51
User-defined Data Types	52
Naming the Table	53

Chapter 3: Keys

Client Specific Performance Objectives and Examples	1-2
Primary Key	3
Foreign Keys	4
Natural and Surrogate Keys	5-6
Natural and Surrogate Keys: Differentiating	7
Referential Integrity	8-9
Primary Key Characteristics	10
Primary Key Constraint	11-12
Composite Key	13

Chapter 4: Design Topics with Normalization

Client Specific Performance Objectives and Examples	1-2
Database Design	3
DB2 Performance Management: Approach	4-5
Normalization	6
Eliminate Repeating Groups - First Normal Form	7-8
Eliminate Redundant Data - Second Normal Form	9-11
Eliminate Columns Not Dependent on Key - Third Normal Form	12-13
Record Size and Page Size	14-16
De-normalization Considerations	17-18
Use Segmented or Partitioned Table Spaces	19-22
Data Compression	23-26
Loading Large Tables	27-28
Index Design Considerations	29
Partitioning Indexes	30-31
When to Build Indexes	32
Ideal Index	33

Chapter 5: Creating DB2 Objects: Part 1

Client Specific Performance Objectives and Examples	1-2
Storage Group	3
DB2-managed Data Sets	4
DB2 Storage Groups - Creating	5
SYSDFLT	6
SMS for Managing DB2 Storage Groups	7
Deferring Allocation of DB2-managed Data Sets	8
CREATE STOGROUP	9-10
Volume Processing	11
CREATE STOGROUP Example	12
Design Implementation	13
Implementing Databases	14
Table Space - Implementation	15
Table Space - Creating Explicitly	16
Table Space - Creating Implicitly	17
Implicit Table Space	18
Tablespace Name	19
Page Size - Choosing	20-21
Page Sizes - Larger	22

Defines the Attributes of a Column	54-55
NOT NULL	56
Constraint and Primary Key	57
CHECK (check-condition)	58
DEFAULT	59
GENERATED	60
GENERATED - AS IDENTITY	61
Minimum and Maximum Values	62
PRIMARY KEY (column-name,...)	63
UNIQUE(column-name,...)	64
FOREIGN KEY (column-name,...) references-clause	65
REFERENCES table-name (column-name,...)	66
LIKE table-name or view-name	67
Creating a Table with an Identity Column	68
Table Example	69
Table Example - Foreign Key	70
DB2 Data Sharing Group	71
Tablespace Scan Access Path	72
Matching Index Scan	73
Non-Matching Index Scan	74
Nested Loop Join (Method=1)	75
Merge Scan Join (Method=2)	76
Hybrid Join (Method=4)	77

Chapter 7: Physical Design Considerations

Client Specific Performance Objectives and Examples	1-2
Data Types	3-4
Byte Counts By Data Type	5-6
Consistent Data Types	7
VARCHAR	8-10
VARCHAR - Performance Savings	11
VARCHAR Recommendation	12
NULL	13
Handling Nulls in the Application	14
Column Placement	15-17
Key Design	18-19
Surrogate Key - Consider When	20
Volatile Tables	21-22
Derived Columns	23
Aggregation	24-25
Materialized Query Tables	26
Partitioning - Benefits	27
Partitioning Considerations	28
DSSIZE Parameter	29
DB2 Compression	30-31
DB2 Compression - When to Use?	32
Cost of Updates on Compressed Rows	33
When Does the Compression Dictionary Get Built?	34-35
Compression Dictionary - When to Rebuild?	36-37
Page Size	38
Rows Per Page	39-40
Page Size and I/O Cost	41

Index - What is it?.....	42
Sample - Three-level Index Structure.....	43
Indexes - Why Build?.....	44
Index - Uniqueness.....	45
Index - Improve Performance.....	46
Index - Cost.....	47
Index - Maintenance.....	48
Indexable Predicates - Matching.....	49-51
Index Screening Predicates.....	52
List Prefetch.....	53-55
Sort Avoidance.....	56-57
Index-only Access.....	58
Index Design Guidelines.....	59
Indexing Strategy - General.....	60-61
Index Design - Multi-column Index or Multiple Indexes.....	62
Predicate Use.....	63
Index Columns - Filtering.....	64
Uniqueness - Enforcing.....	65
Access Path - Efficiency.....	66
Multiple Index Access.....	67
Single Index Access.....	68-69
Index Maintenance.....	70
Column Sequence in a Multi-column Index.....	71
Index Considerations for Join Operations.....	72
Index Design After the First Table.....	73
Indexes - NOT PADDED and PADDED.....	74
NOT PADDED and PADDED Indexes.....	75-76
Clustering Index.....	77
Clustering Index - What is it?.....	78
Clustering versus a Non-clustering Index.....	79
Clustering Index - Choosing.....	80-82
Clustering and the Impact on Asynchronous I/O.....	83
Clustering and the Impact on Join Access.....	84
Clustering Index - Changing.....	85
Partitioned Tables - Indexing.....	86-87
Partitioned Partitioning Index.....	88-89
Nonpartitioned Indexes.....	90
NPI - Advantages.....	91
NPI - Disadvantages.....	92

Chapter 8: Utilities Considerations

Client Specific Performance Objectives and Examples.....	1-2
Checkpoint Frequency.....	3-5
QMF Considerations.....	6
Running Utilities Concurrently.....	7-8
Index Image Copy.....	9-10
LISTDEF.....	11-12
COPY.....	13-17
LOAD Indexes in Parallel.....	18-21
LOAD Partition Parallelism.....	22-23
LOAD Partitions in Parallel.....	24-27

Chapter 10: Application Design and Indexes

Client Specific Performance Objectives and Examples.....	1-2
RUNSTATS.....	3-5
Optimizer Input.....	6
OPTIMIZE FOR n ROWS.....	7
Access Paths.....	8
Nonsegmented Tablespace Scan.....	9-10
Segmented Tablespace Scan.....	11-12
Matching Index Scan.....	13-14
Non-matching Index Scan.....	15-16
Multiple Index Scan.....	17-18
Clustered versus Non-clustered.....	19
Evaluating Predicates.....	20
Compound Indexes.....	21
Nested Loop Join.....	22-23
Merge Scan Join.....	24-27
Hybrid Join.....	28
PLAN_TABLE.....	29
PLAN_TABLE Columns.....	30-33
Sort Indicators.....	34
PLAN_TABLE Columns.....	35
DSN_STATEMENT_TABLE.....	36
Category 'B' Reasons.....	37
DSN_FUNCTION_TABLE.....	38
Schema.....	39

Chapter 11: DB2 Performance Monitoring

Client Specific Performance Objectives and Examples.....	1-2
Materialized Query Table.....	3-5
Materialized Query Table: Creation.....	6
Varying-length Index Keys Performance Benefits.....	7-8
Backward Index Scan.....	9
Stage 1 and Indexable Predicates.....	10
Mismatched Data Types.....	11-12
Unknown Join Sequence.....	13
Trigger Enhancements.....	14
Cost-based Parallel Sort.....	15
EXPLAIN Enhancements.....	16
DB2 Address Spaces.....	17-18
MSTR Address Space.....	19
DBM1 Address Space.....	20
IRLM Address Space.....	21
DIST Address Space.....	22
SPAS Address Space.....	23
Storage Group.....	24-26
Databases.....	27-29
Tablespaces.....	30-38
Tables.....	39-40
Indexes.....	41-43
Free Space: Pros and Cons.....	44

Statistics.....	28-29
REORG and Discard.....	30-31
REORG - What.....	32-33
z/OS Performance Options.....	34-35
LOAD / UNLOAD Delimited Input / Output.....	36
Delimited File.....	37
Delimited Files - LOAD / UNLOAD.....	38
RUNSTATS Enhancements.....	39
Rebalancing: Overview.....	40
LOAD Utility.....	41-42
Delimited Files.....	43
Availability During Load.....	44
Replacing or Adding Data.....	45
Inserting a Single Row.....	46
Inserting Multiple Rows.....	47
Insert.....	48
Load Utility.....	49
LOAD DATA Output.....	50
LOAD Phases of Execution.....	51-52
LOAD Syntax.....	53-56
LOAD Options.....	57-58
Load Statement that Specifies Field Positions.....	59
Loading Selected Records into Multiple Tables.....	60
Loading Data of Different Data Types.....	61
Loading Data in Delimited File Format.....	62
Concatenating Multiple Input Records.....	63

Chapter 9: Catalog in Database Design

Client Specific Performance Objectives and Examples.....	1-2
DB2 Storage Groups.....	3
Table - Information.....	4
Catalog Information About Aliases.....	5
Columns - Information.....	6
Indexes - Information.....	7
Authorizations - Information.....	8
Foreign Keys - Information.....	9
Check-pending - Information.....	10
Check Constraints - Information.....	11
User-defined Functions and Stored Procedures.....	12
Triggers - Information.....	13
Sequences - Information.....	14
Comments - Adding and Retrieving.....	15-16

Indexes.....	45
Pages.....	46-52
Bufferpools.....	53-58
Access via SQL.....	59
SQL Types.....	60
Memory Usage.....	61-62
DBM1 Memory Usage.....	63-64
Memory Usage.....	65
DBM1 Memory Usage.....	66
DDCS Database.....	67-68
DB2 Connect Operation.....	69
RLST Database.....	70-72
Work File Database.....	73-74
Work File Database Usage.....	75-76
TEMP Database.....	77-78
Important Datasets.....	79-80
Parallel Sysplex.....	81
Data Sharing Group.....	82-83
DB2 Accounting Time.....	84

Chapter 12: ERWin

Client Specific Performance Objectives and Examples.....	1-2
Concepts.....	3
Model Types.....	4
Model: Advantages.....	5
Modeling Activities.....	6
Entity - What is it.....	7
Entity Examples.....	8
Relationship - What is it.....	9
Relationship Clarified.....	10
Relationship Types.....	11
Attribute - What is it.....	12
Entity Identifier.....	13
Normalization.....	14
Independent and Dependent Entities.....	15
Attributes.....	16
Primary Key and Non-Keys.....	17
Foreign Key.....	18
Relationships.....	19
Model - Creation.....	20
Blank Model.....	21
Name the Model.....	22
Entities - Adding.....	23
Entity Name - Edit.....	24
Attributes - Adding.....	25
Relationships - Creation.....	26
Text Block - Adding.....	27
Box - Resizing.....	28
Drawing Objects - Adding.....	29
Data Model - Printing.....	30
Headers and Footers - Customizing.....	31

Workplace	32
Main Toolbar	33
Font and Color Toolbar	34
Drawing Objects Toolbar	35
Alignment Toolbar	36
AllFusion Model Manager Toolbar	37
ERwin Toolbox	38
Workplace Panes	39
Model Explorer	40
Data Models	41
Domains	42
Domain Dictionary	43-44
Domain Properties - Inherited and Non-inherited	45
Relationships	46
Entity Relationship Data Model	47
Relationship Types	48
Relationship Lines	49
Relationship Editor	50
Verb Phrases	51
Views	52-53
View to a Model - Adding	54
Physical Display Levels	55
Indexes - Types	56
Index Editor	57
Index - Creation	58
Index Creation - Model Explorer	59
Forward Engineering	60

Chapter 13: Materialized Query Tables

Objectives Client Specific Performance Objectives and Examples	1-2
Materialized Query Tables	3
Automatic Query Rewrite	4
MQT Types	5-8
Summary Tables	9-10
Steps for Using Automatic Query Rewrite	11
Defining a Materialized Query Table	12
Creating a New Materialized Query Table	13
DATA INITIALLY DEFERRED	14
MAINTAINED BY	15
ENABLE QUERY OPTIMIZATION	16
REFRESH TABLE	17-18
INSERT, UPDATE, DELETE, and LOAD for MQT	19
Exploiting Automatic Query Rewrite	20-21
Materialized Query Table Design	22-23

Chapter 15: Sequences and Identity

Client Specific Performance Objectives and Examples	1-2
Common Table Expressions	3
Recursive SQL	4
Controlling Depth of Recursion	5
Identity Columns	6
Identity Column Considerations	7-8
Identity Column Enhancements	9-10
Sequence Objects	11-12
SQL to Support Sequence Objects	13-14
Scalar Fullselect	15
Extension to Expressions Syntax	16
Scalar Fullselect Restrictions	17
Multiple Distinct Clauses	18-20
INSERT within SELECT Statement	21
INSERT within SELECT	22-23
IS NOT DISTINCT FROM	24-25
ROWID Columns as Keys	26
ROWID Access	27
Retrieved ROWID	28
ROWID in Key	29
ROWID Column - Defining	30
Direct Row Access	31
Identity Columns as Keys	32
Identity Column - Defining	33-34

Chapter 14: Security

Objectives Client Specific Performance Objectives and Examples	1-2
Security versus Auditing	3
Controlling Data Access	4
Identifiers	5
Access Control within DB2	6
Administrative Authorities	7
Ownership Privileges	8
Controlling Access to a DB2 Subsystem	9
Data Set Protection	10
Encryption	11
Controlling Access to DB2 Objects	12
Authorization IDs	13
Explicit Database Privileges	14-16
Explicit Package Privileges	17
Explicit Plan Privileges	18
Explicit Table and View Privileges	19-20
Administrative Authorities	21-22
Field-level Access Control by Views	23
Ownership of Objects - Establishing	24
Implicit Privileges - Granting	25
Authorization to Execute	26
Caching Authorization IDs for Best Performance	27
Caching IDs for Packages	28
Recommendation Against Use of Controls	29
User-defined Functions	30
Stored Procedures	31
Sequences	32
List the Views	33
DB2 Authorization IDs with Granted Privileges	34
Retrieving All IDs with DBADM Authority	35
Retrieving IDs Authorized to Access a Table	36
Retrieving All IDs that Can Change the Sample Employee Table	37
IDs that Have Been Granted the Privilege of Updating	38
Tables an ID is Authorized to Access	39
Plans and Packages that Access a Table	40
Multilevel Security	41
Multilevel Security - Advantages	42
Users and Objects	43
Security Labels	44
Implementing Multilevel Security with DB2	45
Row Level Security	46
Working with Data in a Multilevel-secure Environment	47
GRANT Database Privileges	48
GRANT Package Privileges	49
GRANT System Privileges	50
GRANT Table or View Privileges	51