

DB2 UDB® Database Design and Administration

Course No.	1093
Description	This course uses a combination of lecture and hands-on workshops to enable participants to create and administer objects in a DB2 UDB® for Z/OS or OS/390 Version 7 environment. The course also examines how to execute all UDB's utilities including LOAD, CHECK DATA, COPY, QUIESCE, REORG, RECOVER, RUNSTATS, and UNLOAD.
Audience	This course is recommended for Database Administrators, Systems Administrators or Application Developers who are responsible for creating, modifying, or maintaining DB2 UDB® objects.
Prerequisites	A working knowledge of relational database objects is required.
Objectives	<ul style="list-style-type: none">• Use the DB2 UDB® catalog to obtain information about database objects• Perform the physical design of a database• Preserve referential integrity and check constraints• Use proper indexing• Select the correct tablespace types• Create all DB2 UDB® objects• Alter certain DB2 UDB® object characteristics• Understand basic DB2 UDB® performance considerations• Administer a DB2 UDB® database using all DB2 UDB® utilities and Operator commands• Load tables using the LOAD utility• Develop and implement a backup/recovery strategy• Reorganize a tablespace using several different strategies• Unload data from DB2 UDB®• Process multiple objects effectively
Major Topics	<ul style="list-style-type: none">• DB2 UDB® Overview• SQL Overview• The UDB Catalog• DB2 UDB® Table Design and Creation• DB2 UDB® Index Design and Creation• DB2 UDB® Referential Integrity• DB2 UDB® Views, Synonyms, and Aliases• DB2 UDB® Storegroups, Databases, and Tablespaces• DB2 UDB® Security• DB2 UDB® Triggers• DB2 UDB® Locking and Performance• The DB2 UDB® Command Interface• DB2 UDB® LOAD Utility• DB2 UDB® Backup and Recovery• DB2 UDB® REORG and UNLOAD• Multi-object Processing• Standalone Utilities
Duration	5 days



Course Contents

1. DB2 UDB® Overview

- What DB2 UDB® Is
- How DB2 Interfaces with Applications
- Components of DB2 UDB® Database
- DB2 UDB® “World” Example
- String Data Types
- Binary Data Types
- Numeric Data Types
- Date/Time Data Types
- Attribute Qualifiers
- Check Constraints
- DB2 UDB® “World” Example
- Relationships
- A View
- Views
- ‘Special’ Data Types
- Table/View names
- Synonyms and Aliases A DB2 UDB® Sybssystem

2. SQL Overview

- SQL
- SQL Syntax Rules
- DDL, DCL, DML
- Static SQL
- Dynamic SQL
- Mainframe Dynamic SQL
- Workstation Access – The Control Center

3. The DB2 UDB® Catalog

- DB2 UDB® System Tables
- Contents of Catalog
- Database: SYSDATABASE
- Tables: SYSDATABASE Columns: SYSCOLUMNS
- Indexes: SYSINDEXES
- Tablespaces and Partitions: SYSTABLESPACE and SYSTABLEPART
- Stored Procedures: SYSROUTINES
- Trigger: SYSTRIGGERS
- Authorizations: Overview
- Resource Authorizations
- System-Wide Authorizations
- Database Authorizations
- Table Authorizations
- Column Authorizations

- Package Authorizations
- Plan Authorizations
- DB2 UDB® Directory Catalog Query Tools

4. Table Design and Creation

- Logical Design
- Physical Design
- A Simple Data Model
- Transforming Data Model to Physical Design
- Transforming Many-to-Many Relationships
- Transforming Complex Relationships
- Normalization
- Intentional Denormalization
- Denormalization Techniques
- Identifying Denormalization Opportunities
- Table Creation Process Naming the Table
- Defining the Columns
- Numeric Data Types
- Character/Graphic Data Types
- Date/Time Data Types
- LOB Data Types
- Special Data Types
- The ROWID Datatype
- The AS IDENTITY Datatype Modifier
- Null/Default Characteristic
- Placing A Table In A Database
- Table Limits
- Tables As Related To Tablespaces Column Constraints
- Constraints At The Table Level
- Column Constraints
- Deleting a Table with DROP
- DROP/ReCREATEing A Table
- Restricting DROP
- ALTERing A Table
- Sample ALTER Commands
- Creating a Table Like Another
- Renaming a Table
- Global Temporary Tables

5. Index Design and Creation

- Index Benefits
- Index Costs
- Required Indexes
- Optional Indexes



- CREATE Index DDL
- CREATE INDEX Parameters
- Index Structure
- Altering Index Characteristics
- Dropping Indexes
- Drop/ReCREATE INDEX
- Estimating Space for Indexes
- Non-Unique Index Space Estimation

6. Referential Integrity

- Definition of R.I.
- R.I. Rules
- Delete Rules
- Notes on R.I.
- Creating Tables and Indexes with RI
- FK Indexes
- Implementing R.I. with DDL
- CREATE vs. ALTER for R.I.
- Recursive relationships
- Cyclic Relationships
- Application-Enforced R.I.

7. Views, Synonyms & Aliases

- Definitions
- Why Views?
- Column Subset Views
- Row Subset Views
- Changing Columns Names Via Views
- Hiding Complex SQL with View
- Updatability of Views
- Views WITH CHECK OPTION
- Views WITH CASCADED CHECK OPTION
- Synonyms and Aliases
- CREATE Synonym
- CREATE Alias

8. Databases, Tablespaces And Storage Groups

- Definitions
- Storage Groups
- Databases
- Databases and Bufferpools
- Databases and Characters Sets
- CREATE DATABASE
- Creating Tablespaces
- CREATE Simple/Segmented Tablespace DDL
- Tablespace Types

- Simple vs. Segmented Tablespaces
- Segmented Tablespaces
- Multiple Tables per Tablespace
- Implicit vs. Explicit Tablespace Creation
- CREATE Tablespace Parameters
- Dataset Naming Convention
- Altering Simple Segmented Tablespaces
- Partitioning Your Data
- Partitioned Tablespace Advantages
- Partitioning Drawbacks
- Partitioned Tablespaces
- 'Normal' Partitioned Tablespaces
- Customizing Partitions
- CREATE 'Normal' Partitioned Tablespace DDL
- 'Normal' Partitioned Tablespace Sample DDL
- Finishing the 'Normal' Partition CREATE Process
- The Partitioning Index
- Indexes on Partitioned Data
- Partitioning and Index Datasets
- Rebalancing Data Within Partitions
- Gigantic Partitioned Tablespaces
- Gigantic Partitioned Tablespace DDL
- Managing LOB Data
- Estimating Space for Tables and Tablespaces
- Physical structures
- Manually Managing VSAM Files (VCAT)

9. Security

- Obtaining Your UDB Identity
- AUTHIDs
- Sources of Authorization
- GRANT DCL
- REVOKE DCL
- System Privileges
- Resource Privileges
- Databases Privileges
- Table Privileges
- Plan/Package Privileges
- Other Privileges
- Administrative Groups
- System Administrative Groups
- Database Administrative Groups

10. DB2™ Triggers, User Defined Function & User defined Types



- What's a Trigger?
- Why Do We need Triggers?
- Some Common Uses of Triggers
- Trigger Definition – Name, Activation Time, Triggering Event, Base Table, Granularity
- Activation Time and Granularity
- Trigger Definition – Mode
- Triggers: Transition Variables & Tables
- Referencing Transition Variables
- Trigger Definition – Conditions, Action
- Calling Procedures from Triggers
- Functions from Triggers Unconditionally
- Calling User Defined Functions from Triggers Conditionally
- Signaling Error Conditions from Triggers
- Dropping a Trigger
- Order of Multiple Triggers
- Trigger Cascading
- Trigger and Referential Constraints
- User Defined Functions
- Types of UDFs
- Creating UDFs
- UDFs Shipped with DB2
- DAYNAME
- MONTHNAME
- UDT

11. Expanding Your DB2Horizons

- BLOBs
- CLOBs
- DBCLOBs
- LOB Locators
- Using LOBs in A Program
- Using LOB Locators in a Program
- LOB Functions
- Other Large Object Facts
- Stored Procedures
- Stored Procedures in the Catalog
- Stored Procedure Considerations
- Dynamic SQL in Programs
- Data Sharing
- Batch SQL
- DSNTIAUL Example

12. DB2™ Locking and Performance

- Locking

- SQL Performance Factors
- Access Paths and EXPLAIN
- Catalog Tables Updated by RUNSTATS
- Access Paths and EXPLAIN
- EXPLAIN
- Access Paths and EXPLAIN
- Taking Control With Optimization

13. UDB® Utilities Overview

- Common DB2™ Utilities
- DB2™ Objects
- Common DB2 Commands
- Displaying Objects
 - DISPLAY Database
- Other Options
- DB2I Option 7
- Status Codes
 - START Database
 - STOP Database
- Logical/Physical Design for Lab

14. The UDB LOAD Process and Utilities

- Alternatives for Loading Tables
- Application Program Load Algorithm
- Application Program vs. LOAD Utility
- Load Capabilities
- LOAD Overview
- LOAD JCL
- LOAD JCL Components
- LOAD Schematic
- Utility Phases
- Load Phases
- LOAD Control Statement
- LOADING a File That Matches the Table Definition
- LOADING a File That Does Not Match the Table definition
- LOAD Statement Notes
- COBOL Datatypes vs. Load Datatypes
- The DB2I Utility Panel
- Utility IDs
- Tracking Utilities
- LOAD Return Codes/Abends
- Logical/Physical Design for Lab
- Lab DDL – Customers Table
- Load Control Statement Options
- LOAD – Append vs. Overlay



- High Performance LOAD – Logging
- High Performance LOAD – R.I.
- LOAD Discard Limits
- LOADING Partitions In Parallel
- LOAD – Compressed Data
- LOADING Only Selected Input Records
- Massaging Data As It Is Loaded
- Combining LOAD Parameters
- LOADING More Than One Table
- Overriding Default DD Names
- High Performance LOAD – Parallel Index Build
- Creating an Image Copy When LOADING
- Preformatting Pages with LOAD

15. The UDB® LOAD Process and Utilities

- LOADING Data Extracted by REORG
- Gathering Statistics When LOADING
- v7 LOAD Enhancements
- LOAD Performance Tips
- Gathering Statistics Via RUNSTATS Statements
- PUNSTATS Parameters
- Removing Historical Statistics
- RUNSTATS Correlated Statistics
- Removing “Pending” Statuses after LOADING
- “Parent” CHKP Events
- “Child” CHKP Events
- Removing CHKP Status
- The Check Data Utility
- Removing Erroneous Data Using the Check Data Utility

16. UDB® COPY and RECOVER Utilities

- Summary of Recover-Related Utilities
- Important Recovery Factors
- Recovery and R.I.
- Developing a Backup/Recover Strategy
- Your Backup Strategy
- Effective Backup/Recovery
- The COPY Utility
- Specifying Objects to COPY
- COPY Parameters – DD Names
- COPY Parameters – Performance
- COPY Parameter – Automating the COPY Type Decision
- Copying a List of Objects
- COPY – Other Notes
- Image Copies in SYSCOPY

- Recovery Enhancement Utilities
- The QUIESCE Utility
- QUIESCE Parameters
- The REPORT Utility
- The Modify RECOVERY Utility
- The Modify Utilities
- The RECOVER Utility
- Point-In-Time Recovery
- What Makes Recovery Necessary?
- The RECOVER Utility
- Eight Steps to Recovery
- RECOVER Parameters
- RECOVER Parameters – Performance
- Controlling the Recovery Point
- Tablespace Recovery Example
- Notes on RECOVER
- Rebuilding An Index Using REBUILD INDEX
- Notes on Point-In-Time Recovery
- DSN1LOGP Notes
- The COPYTOCOPY Utility
- The MERGECOPY Utility
- Stacking Image Copies on Tape

17. UDB® REORG and UNLOAD Utilities

- The REORG Utility
- Reasons to REORG
- Off-Line vs. Online REORG
- Statistics That May Drive REORG
- REORG Schematic
- Off-line REORG Phases
- REORG Datasets
- REORG Parameters – Type of REORG
- REORG Parameters – Performance
- REORG Parameters – Unloading Only
- REORG Parameters – Autoinitiation
- RESTART
- Return Codes/Abends and Restartability
- Online REORG Phases
- Online REORG
- Online REORG – REFERENCE Considerations
- Online Considerations
- Online REORG Mapping Table
- REORG INDEX
- REORG INDEX – Parameters



- CHECK INDEX
- The UNLOAD Utility
- UNLOAD Control Statement
- UNLOAD Parameters

18. Multiple Object/Dataset Processing

- LISTDEF's Purpose
- Utilities That Support LISTDEF
- LISTDEF Syntax
- LISTDEF Processing
- Using LISTs With Utilities
- TEMPLATE's Purpose
- Utilities That Support TEMPLATE
- TEMPLATE Syntax

- TEMPLATE Variables
- Other TEMPLATE Options
- OPTIONS's Purpose
- OPTIONS Syntax

19. UDB® Standalone Utilities

- DSN1COMP – Estimate Compression
- DSN1COPY – COPY A VSAM File
- DSN1LOGP – Print The Recovery Log
- DSN1LOGP Notes
- DSN1PRNT – Print A DB2 Data File

Appendices

- A. Limit in DB2™ for OS/390

