



TERADATA EDUCATION OUTLINE

Coffing Data Warehousing has provided quality Teradata education, products and services for over a decade. We offer customized solutions to maximize your warehouse.

Toll Free: 1-877-TERADAT
Business Phone: 1-937-855-4838
Email: <mailto:CDWSales@CoffingDW.com>
Website: <http://www.CoffingDW.com>

In addition to the course material listed in this outline, we also offer Teradata classes in Teradata Basics, Implementation, SQL, Database Administration, Design and Utilities. Please contact us so we can customize a course to fit your specific needs.

PURPOSE

Coffing Data Warehousing has been providing quality Teradata education for over a decade. We offer customized courses to maximize the effectiveness of each class. The purpose of this proposal is to build a lasting relationship with your company. To this end, we have combined our comprehensive Teradata education services in a unique package that we feel best suits the diverse needs of your company while offering our high quality product at competitive pricing.

Coffing Data Warehousing is excited to offer you, our preferred partner, an innovative new way to look at training at the CoffingDW Teradata University (CDW-TU). This approach provides the ability to maximize learning potential. Our goal is to make your employees the most educated data warehouse experts in the industry.

CURRICULUM:

Coffing Data Warehousing will provide an experienced and highly qualified resource to deliver this customized educational seminar on the following topic(s):

Teradata Education

• Teradata Application Development

COURSE DESCRIPTION

COURSE PREREQUISITES	There is no prerequisite for this course.
COURSE DURATION/FORMAT	This course is designed to be highly interactive with the audience.
COURSE AUDIENCE	The audience will consist of a mix of beginning, intermediate and advanced Teradata users.
OBJECTIVES	This course is designed to provide in-depth knowledge of Teradata Application Development.

Tera-Tom on Teradata Application Development for V2R6

Chapter 1 — The Rules of Data Warehousing

Teradata Certification

A Logical View of the Teradata Architecture

The Parsing Engine (PE)

The Access Module Processors (AMPs)

The BYNET

A Visual for Data Layout

How Teradata handles data access

Teradata Cabinets, Nodes, Vprocs, and Disks

Three Types of Parallelism with Teradata

Multi-Step Parallelism

Multi-Statement Parallelism

The Active Data Warehouse

OLTP Environments

The DSS environment

Teradata Maximums

Teradata Maximums per Release

Chapter 2 — Data Layout: AutoMagically

Data Distribution Explained

Rows and Columns

The Primary Index

The Two Types of Primary Indexes

Unique Primary Index (UPI)

Non-Unique Primary Index

Data Layout Summary

Chapter 3 — Partition Primary Indexes

V2R5 Partition Primary Indexes

V2R4 Example

V2R5 Partitioning

Partitioning doesn't have to be part of the Primary Index

Partition Elimination can avoid Full Table Scans

The Bad NEWS about Partitioning on a column that is not part of the Primary Index

Two ways to handle Partitioning on a column that is not part of the Primary Index

Partitioning with CASE_N

Partitioning with RANGE_N

NO CASE, NO RANGE, or UNKNOWN

Chapter 4 — The Extended Logical Data Model

The Application Development Life Cycle

Asking the Right Questions

Logical Data Model

Primary Keys

Foreign Keys

Normalization

A Normalized Data Warehouse

Dimensional Modeling

Extended Logical Data Model

The End Goal of the ELDM is to build Table Templates

Column ACCESS in the WHERE Clause

Data Demographics

Extended Logical Data Model Template

The Physical Data Model

How to Pick an excellent Primary Index

Denormalization

Derived Data

Temporary Tables

Derived Tables

Volatile Temporary Tables

Global Temporary Tables

Chapter 5 — Secondary Indexes

Unique Secondary Index (USI)

USI Subtable Example

How Teradata retrieves an USI query

NUSI Subtable Example

How Teradata retrieves a NUSI query
Value-Ordered NUSI
Collecting Statistics on NUSI Indexes
NUSI Bitmapping
Prototyping indexes with EXPLAIN
Chart for Primary and Secondary Access
Secondary Index Summary

Chapter 6 — Join Strategies

A Join in Simple Terms
The key things to know about Teradata and Joins
Merge Join Strategies
Merge Join Strategies
Merge Join Strategies
Joins need the joined rows to be on the same AMP
Another Great Join Picture
Joining Tables with matching rows on different AMPs
Redistributing a Table for Join Purposes
Big Table Small Table Join Strategy
Big Table Small Table Duplication
Nested Join
Hash Join
Exclusion Join
Product Joins
Cartesian Product Join
Outer Joins
Inner and Outer tables with Residual Conditions

Chapter 7 — Join and Hash Indexes

Description of Join Indexes
Join Index Fundamentals
Types of Join Indexes
Single-Table Join Indexes
Multi-table Join Indexes
Aggregate Join Index
Partial-Covering Global Join Index
Partial-Covering Global Join Index Picture

Partial-Covering Global Join Index –
Multi-Table Join Back
Sparse Index
Sparse Index Picture
UPSERT ON Indexed Tables
Hash Indexes
Hash Indexes vs. Secondary Indexes
Hash Indexes vs. Single-Table Join Indexes
Hash and Join Indexes vs. Base Tables

Chapter 8 — Explains

The Teradata Optimizer “knows” how to Explain
Row Estimate Confidence Levels
Explain Terminology

Chapter 9 — Understanding Views, Macros and Triggers

All About Views
View Aggregation
All About Macros
Creating a MACRO
Macros that Use Parameters
All About Triggers

Chapter 10 - Locks

Teradata has 4 locks for 3 levels of Locking
Locks and their compatibility
Locks and their compatibility
How Teradata Locks Objects
Teradata Locks – First Come First Serve
Locking Queue Example 2
Locking Queue Example 3
Locking Modifier
The NOWAIT Option

Chapter 11 - Collect Statistics

How Collect Statistics Works
Sample Statistics
Sample Statistics
What You Should COLLECT STATISTICS On

Chapter 12 — Tools

Database Query Log (DBQL)
DBQL Collection Options
DBQL Tables and Views
How to Begin Logging for DBQM
How to Begin Logging for DBQM
Performance Monitor - PM
Access Logging
Statistics Wizard
Index Wizard
TSET

Chapter 13 — Loading the Data

Fastload
Multiload
Tpump
Fastexport
Insert/Select
Insert/Select

Chapter 14 — MISC

Identity Columns
Identity Column Example
Set Manipulation
System Calendar
Stored Procedures
OLAP and Analytical Functions
ANSI OLAP Syntax
Referential Integrity
Soft Referential Integrity
Materialized Views

Roles

Profiles

Compression

Implementing Compression

How Compression Works

Teradata and ANSI Mode

Teradata Multi-Statement Request Examples

PP2 –The Teradata Preprocessor

Call Level Interface - CLI