



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 Designer**

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 Designer.

Tera-Tom on Teradata Designer 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
Teradata Cabinets, Nodes, Vprocs, and Disks
Teradata Maximums
Teradata Maximums per Release

Chapter 2 — Data Distribution Explained

Rows and Columns
The Primary Index
The Two Types of Primary Indexes
Unique Primary Index (UPI)
Non-Unique Primary Index
Turning the Primary Index Value into the Row Hash
The Row Hash Value determines the Rows Destination
The Row is Delivered to the Proper AMP
The AMP will add a Uniqueness Value
Teradata accesses data in three ways
Data Layout Summary

Chapter 3 — 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
Partitioning and Joins

Chapter 4 — Teradata Space

How Permanent Space is calculated
How Permanent Space is Given
The Teradata Hierarchy
How Spool Space is calculated
A Spool Space Example
PERM, SPOOL and TEMP Space
Spool Space controls system time
Spool Reserve
A quiz on Perm and Spool Space
Another quiz on Perm and Spool Space

Chapter 5 — Modeling the Data

The Application Development Life Cycle
Asking the Right Questions
Logical Data Model
Primary Keys
Foreign Keys
Normalization
A Normalized Data Warehouse
Dimensional Modeling
Data Marts
Update Applications

Chapter 6 — Extended Logical Data Model

The End Goal of the ELDM is to build Table Templates
Column ACCESS in the WHERE Clause
Data Demographics
Distinct Values
Typical Rows Per Value
Maximum Rows NULL
Change Rating

Extended Logical Data Model Template

Chapter 7 — Denormalization

- Derived Data
- Horizontal Partitioning
- Vertical Partitioning
- Multi-Table Join Indexes
- Temporary Tables
- Derived Tables
- Volatile Temporary Tables
- Global Temporary Tables
- Views

Chapter 8 — 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
- Value-Ordered NUSI
- How Teradata retrieves a Value-Ordered NUSI query
- NUSI Bitmapping
- Prototyping indexes with EXPLAIN
- Chart for Primary and Secondary Access
- Secondary Index Summary

Chapter 9 — Join Strategies

- A Join in Simple Terms
- The key things to know about Teradata and Joins
- 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

Chapter 10 — Join Indexes

Three basic types of Join Indexes
Join Index Fundamentals
Join Indexes versus other objects
Multi-Table Join Index
Single-Table Join Indexes
Aggregate Join Index
Sparse Index
Sparse Index Picture
Global Join Index
Global Join Index Picture
Global Join Index – Multi-Table Join Back
Hash Indexes
Hash Indexes vs. Single-Table Join Indexes

Chapter 11 - Locks

Teradata has 4 locks for 3 levels of Locking
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 12 — Tools

Database Query Log (DBQL)
DBQL Collection Options
DBQL Tables and Views

How to Begin Logging for DBQM
Teradata Dynamic Query Manager (TDQM)
TDQM works with Queryman (SQL Assistant)
TDQM Workloads and Rules
Performance Monitor - PM
Access Logging
Statistics Wizard
Index Wizard
TSET

Chapter 13 — Loading the Data

Fastload
Fastload Picture
Multiload
Multiload Picture
Tpump
Tpump Picture
Fastexport
Fastexport Picture
Warehouse Builder
Insert/Select
Utilities Summary

Chapter 14 — MISC

Identity Columns
Identity Columns Example
Referential Integrity
Soft Referential Integrity
Roles
Profiles
Compression
Implementing Compression
How Compression Works
Sync Scan
Call Level Interface - CLI
LAN Connection to Teradata for Network Attached Clients
Mainframe Connection to Teradata

Chapter 15 — Priority Scheduler

Priority Scheduler Partition Hierarchy
Priority Scheduler Hierarchy Definitions
Resource Partition Example
Multiple Resource Partitions Example
Scheduling Policies
Performance Periods

Chapter 16 — Data Protection

Transaction Concept & Transient Journal
How the Transient Journal Works
FALLBACK Protection
How Fallback Works
Fallback Clusters
Down AMP Recovery Journal (DARJ)
Redundant Array of Independent Disks (RAID)
Cliques
Cliques – A two node example
Cliques – A four node example
Permanent Journal
Table create with Fallback and Permanent Journaling
Archive and Recovery
Disaster Recovery
Teradata Crash Dumps