



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 Physical Implementation

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 Physical Implementation.

Tera-Tom on Teradata Physical Implementation 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

Teradata Cabinets, Nodes, Vprocs, and Disks

Chapter 2 — Data Distribution Explained

Rows and Columns

The Primary Index

The Two Types of Primary Indexes

Unique Primary Index (UPI)

Turning the Primary Index Value into the Row Hash

The Row Hash Value determines the Row's Destination

The Row is Delivered to the Proper AMP

The AMP will add a Uniqueness Value

An Example of an UPI Table

An Example of a NUPI Table

How Teradata Retrieves Rows

Row Distribution

A Visual for Data Layout

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 Engine – Under the Hood

Full Cylinder Read
Table Header
Each Table is given a Table ID
How Data Blocks are Dynamically Built
Data Blocks
How Teradata Finds a Row of Data
The Master Index
The Cylinder Index
Cylinder Index Changes
How Teradata Writes to an AMP
Writing to Data Blocks of Equal Length
When a Data Block is not Big Enough for a Write
How Teradata Allocates Blocks
Block and Row Definitions
Large Row versus Oversized Row
Defragmentation
When a Cylinder becomes Full
A Node and its Memory Allocations

Chapter 5 — The Extended Logical Data Model

The Application Development Life Cycle
Asking the Right Questions
Logical Data Model
Primary Keys
Foreign Keys
Normalization
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 6 — The Physical Data Model

Step 1 – Look at Distribution
Step 2 – Eliminate based on Change Rating
Step 3 – NUSI Elimination via Value Access Frequency
Step 4 – Pick the Primary Index
Primary Index Factors
Why Join Access Frequency is Top Priority?
Why Value Access Frequency is Second Priority?
What have we learned about picking the Primary Index?
Step 5 — Pick Secondary Indexes
USI to eliminate Duplicate Row Checking
NUSI considerations
Multi-Column NUSI Columns used as a Covered Query
Value-Ordered NUSIs
A Strongly Vs Weakly Selective NUSI
A formula for calculating a strongly selective NUSI
Typical Row Size
Typical Block Size

Chapter 7 — Denormalization

Derived Data
Storing Aggregates
Pre-Joining Tables
Repeating Groups
Horizontal Partitioning
Vertical Partitioning
Covered Query
Single-Table Join Indexes
Multi-Table Join Indexes
Temporary Tables

Derived Tables
Volatile Temporary Tables
Global Temporary Tables

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
How Teradata retrieves a Value-Ordered NUSI query
NUSI Bitmapping
Prototyping indexes with EXPLAIN
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
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 — Explains

- The Teradata Optimizer “knows” how to Explain in Detail
- Row Estimate Confidence Levels
- Explain Terminology
- Visual Explain

Chapter 12 — The Parsing Engine in Detail

- The Parsing Engine (PE) goes through Six Steps
- Each PE has a Plan Library called RTS Cache
- The Parsing Engine has Data Dictionary Cache
- Why the PE loves the Macro
- The Parsing Engine in Detail
- The Parsing Engine Knows All

Chapter 13 — Understanding Views and Macros

- Views
- Creating Simple VIEWS and VIEWS that Join Tables
- How to Change a VIEW Using REPLACE
- How to Drop a VIEW
- View Aggregation and Nested Views
- All About Macros
- Macros that Use Parameters
- Changing a MACRO Using REPLACE
- How to Drop a MACRO

Chapter 14 - Locks

Teradata has Four locks at Three levels
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 15 - Collect Statistics

Dynamic AMP Sampling
How Collect Statistics Works
Sample Statistics
Sample Statistics
What You Should COLLECT STATISTICS On
COLLECT STATISTICS DETAILED SYNTAX
COLLECT STATISTICS Examples

Chapter 16 - MISC

Identity Columns
Identity Columns Example
LIKE Clause
SUBSTRING and SUBSTR Functions
Referential Integrity
Soft Referential Integrity
Materialized Views
Compression
Implementing Compression
How Compression Works
Creating a Table With DATABLOCKSIZE
The ALTER TABLE Command
Index Wizard
Index Wizard

