



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 Basics**

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 Basics.

Tera-Tom on Teradata Basics for V2R6

Chapter 1 — The Rules of Data Warehousing

Teradata Facts and Certification
Teradata: Brilliant by Design
The Teradata Parallel Architecture
A Logical View of the Teradata Architecture
The Parsing Engine (PE)
The Parsing Engine in Detail
The Parsing Engine Knows All
The Access Module Processors (AMPs)
The BYNET
A Visual for Data Layout
Teradata is a shared nothing Architecture
Teradata has Linear Scalability
How Teradata handles data access
The PE uses Statistics to come up with the Plan
When there are NO Statistics Collected on a Table
Teradata Cabinets, Nodes, VPROCs, and Disks
A Node and its Memory Allocations
LAN Connection for Network Attached Clients
Mainframe Connection to Teradata

Chapter 2 — Data Distribution Explained

Rows and Columns
The Primary Index
Two Types of Primary Indexes (UPI or NUPI)
Unique Primary Index (UPI)
Non-Unique Primary Index
How Teradata Turns the Primary Index Value into the Row Hash
The Row Hash determines the Rows 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 an NUPI Table
How Teradata Retrieves Rows with the Primary Index

Row Distribution
A Visual for Data Layout
Teradata accesses data in three ways
Data Layout Summary

Chapter 3 — 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
A quiz on Perm and Spool Space
Another quiz on Perm and Spool Space

Chapter 4 — 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 5 — 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

Locks

Teradata has 4 locks for 3 levels of Locking

Locks and their compatibility

Chapter 6 — Loading the Data

FastLoad

FastLoad Picture

Multiload

Multiload Picture

TPump

TPump Picture

FastExport

FastExport Picture

Chapter 7 — 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

Secondary Index Summary

Chart for Primary and Secondary Access

Chapter 8 — The Active Data Warehouse

OLTP Environments

The DSS environment

Mixing OLTP and DSS environments

Detail Data

Easy System Administration
Data Marts
Logical and Dependent Data Marts
Teradata Tools - SQL Assistant
TDQM
Index Wizard
Archive Recovery
Teradata Analyst Suite