



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 Database Administration**

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 Database Administration.

Tera-Tom on Teradata DBA Tera-Cram 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 Request and Respond Parcel
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
Each PE has a Plan Library called RTS Cache
LAN Connection for Network Attached Clients
Mainframe Connection to Teradata
Sessions and Session Pools
Teradata Configuration Utilities
Config and Reconfig

Chapter 2 — 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
Space Allocation Review
AMP Disks have Cylinders and Data Blocks

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
Another quiz on Perm and Spool Space

Chapter 3 — DBC Data Dictionary Tables

Data Dictionary Directory
The Parsing Engine has Data Dictionary Cache
Data Dictionary Directory Tables
System Views
Accessing Restricted Views
Selecting Information about Created Objects
Children View
Databases View
DBC.Users View
Indices View
AllTempTables View
Finding Table Names Using the LIKE Command
Finding Table Names in a Particular Database
Using the Keyword USER on DBC Views
Using DBC.AMPUsage
Using DBC.TableSize
Keeping Track of Logons and Logoffs

Chapter 4 — Access Rights, Roles and Profiles

Access Rights

Automatic, Implicit, and Explicit Rights

Tools for Finding Access Rights

The REVOKE Statement

Roles

Creating Roles

Setting and MODIFYING Roles

DBC.RoleInfo and DBC.RoleInfoX

Profiles

CREATING PROFILES

MODIFYING PROFILES

DBC.ProfileInfo and DBC.ProfileInfoX

Chapter 5 — DBS Control

DBS Control Record — General Fields

DBS Control Record — File System Fields

DBS Control Record — Performance Fields

DBS Control Records You Should Know About

DBSControl — Performance Columns

Chapter 6 — Query Analysis and Tools

Database Query Log (DBQL)

DBQL Collection Options

DBQL Tables and Views

How to Begin Logging for DBQM

Access Logging

Statistics Wizard

Index Wizard

TSET

Teradata Visual Explain Utility

Chapter 7 — Teradata Protection Features

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
TDQM

Chapter 8 — Starting and Stopping Teradata

Restarts of the Teradata Database
Automatic Restarts
Hardware Failures
Critical Database Errors
UNIX Operating System Restarts
DBA Forced Restarts
Restarting in UNIX
Restarting the DB Window (UNIX Only)
Restarting in Windows 2000
Startup and Recovery
Transaction Recovery
Performing Online and Offline Catch-up

Chapter 9 — Databases, Users and Accounting

Creating a Database
Creating a User
Modifying and Deleting a USER
Specifying Account Priorities
System Accounting
System Accounting Views
DBC.AccountInfo[x] View
DBC.AMPUsage View

Chapter 10 — Views and MACROS

Creating Views

Creating Simple VIEWS and Views that Join Tables

How to Change a VIEW Using REPLACE

How to Drop a VIEW

View Aggregation

Using “Locking for Access” in Views

You can UPDATE Tables through Views

Restricting UPDATE rows WITH CHECK OPTION

Creating Macros

Why the PE loves the Macro

Creating a MACRO

Macros that Use Parameters

Changing a MACRO Using REPLACE

How to Execute a MACRO

How to Drop a MACRO

Chapter 11 — System Access Control Levels

Teradata Password Encryption

Password Security Features

Host Logon Processing

GRANT/REVOKE LOGON Statement

Session Related Views

DBC.SessionInfo View

DBC

DBC

Data Access Information Views

AccLogRules Views

AccessLog Views

Chapter 12 — Priority Scheduler

Priority Scheduler Partition Hierarchy

Priority Scheduler Hierarchy Definitions

Resource Partition Example

Multiple Resource Partitions Example

Scheduling Policies

Performance Periods

Chapter 13 — Teradata Manager

Teradata Manager Applications
Teradata Dashboard — New Feature
Workload Trend Analysis/Data Collector — New Feature
Priority Scheduler Administrator — New Feature
Priority Scheduler Administrator — New Feature
Teradata Manager Service
Starting Teradata Manager
Enable Data Collection

Chapter 14 — Monitoring Tools

Performance Monitor - Overview
Performance Monitor
Performance Monitor — Continued
PMON Main Window
PMON Sessions Screen
Viewing Session Status
Monitoring Session Status
Session Status Report Descriptions
Teradata Administrator (WinDDI)
Teradata Manager Dynamic Utilization Charting
The Alert Facility and Viewer
The Alert Viewer
The Alert Policy Editor
Locking Logger
Locking Logger Functions:
Xperfstate
Teradata Manager Remote Console

Chapter 15 — Teradata Remote Console Utilities

Starting the Database Window (DBW)
QRYCONFIG
QRYSESSN
RCVMANAGER (Recovery Manager)
SHOWLOCKS

VPROCMANAGER
FERRET UTILITY (File Reconfiguration Tool)
SHOWSPACE
SHOWBLOCKS
PACKDISK
SCANDISK

Chapter 16 — Resource Usage Data

ResUsage Collection and Logging
Setting Resource Logging - DBW
Collection Costs
Resource Usage Tables
ResNode Macros
ResNode Macros - Continued
RSSMon

Chapter 17 — Loading the Data

Fastload
FastLoad Has Some Limits
Three Key Requirements for FastLoad to Run
Maximum of 15 Loads
FastLoad Has Two Phases
PHASE 1: Acquisition
PHASE 2: Application
Fastload Example
Restarting FastLoad
How the CHECKPOINT Option Works
Restarting with CHECKPOINT
MultiLoad
Two Multiload Modes: IMPORT and DELETE
Block and Tackle Approach
MultiLoad Imposes Limits
RESTARTing Multiload
RELEASE MLOAD: When You DON'T Want to Restart Multiload
TPump
Why It Is Called "TPump"
TPump Has Many Unbelievable Abilities

TPump Has Some Limits
LOAD Parameters *IN COMMON* with MultiLoad
.BEGIN LOAD Parameters *UNIQUE to TPump*
A Simple TPump Script — A Look at the Basics
TPump Script with Error Treatment Options
RESTARTing TPump
TPump and MultiLoad Comparison Chart
Fastexport
How FastExport Works
FastExport Fundamentals
FastExport Supported Operating Systems
Maximum of 15 Loads

Chapter 18 — Archiving Data

Archive and Recovery Statements
Recovery Vs FastLoad
Invoking Archive
Invoking Archive - Continued
Restart Log
ANALYZE Statement
Archive
Database DBC Archive
Indexes Option
Database DBC Archive
Archive and Recovery (ARC) Examples

Chapter 19 — Restoring Data

The RESTORE Statement
COPY Statement
Copying Tables
BUILD Statement
RELEASE LOCK Statement