
**Teradata Consulting, Professional Services
and Training**

LEADING EDGE >

Affordable, Experienced, Certified

Teradata Training Courses Overview:

Teradata Utilities

Teradata Physical Implementation/Database Design

Teradata Database Administration

Teradata Application Development

Company Details

Contacts

Terry Graham
Sales Director
+44 (0) 7768 051155
tgraham@leadingedge.uk.com

Richard Cutts
Operations Director
+44 (0) 7711 213101
rcutts@leadingedge.uk.com

Office Address

Egerton House
2 Tower Road
Birkenhead
Wirral
UK
CH41 1FN

Office Telephone

+44 (0) 870 75 76 750

Office Fax

+44 (0) 151 649 9129

Web site

www.leadingedge.uk.com

Teradata Utilities, Teradata Physical Implementation and Database Design,

Teradata Database Administration and Teradata Application Development

Bringing Choice To Teradata Training

The Teradata consultants at Leading Edge have been delivering quality Teradata solutions for over 10 years. In fact, several of our consultants have experience dating back to almost the arrival of Teradata in Europe. We are firmly committed to NCR Certification and boast Certified Teradata Masters.

Through feedback from clients we increasingly recognised the need for an alternative Teradata training offer in the European marketplace.

Rather than 'reinvent the wheel' and develop our own courseware from scratch, we are delighted to have teamed up with Coffing Data Warehousing (CDW) from the US. This strategic alliance brings quality, tried and trusted Teradata courseware to the European Teradata market.

Finally, there is a real choice in the Teradata training market.

What makes our courses special?

We are committed to a teaching style that is interactive, fun, and easy to understand. The important content is emphasised so it becomes extremely clear. Our materials are brilliantly written with all the latest and greatest Teradata content.

We promise not to stand at the front and simply read notes at the students. In our opinion that approach is simply not training!

Why do customers love to deal with us?

We cover the entire Teradata syllabus from Teradata Basics through to the most advanced subject matter. All courses can be tailored to your exact requirements.

We don't publish a schedule in advance – our courses are tailored to your specific requirements, including the timing and the number of students.

All courses are based on the highly acclaimed series of Teradata books published by Coffing Data Warehousing. The examples we work through in class are those in the relevant book. Each student keeps the book(s) used during the class to take away and use for future reference.

Teradata Utilities – Overview

Prerequisites	There are no prerequisites for this course.
Duration	This is a 3-day course designed to be highly interactive with the audience.
Courseware	“Teradata Utilities – Breaking the Barriers”.
Audience	The audience will normally consist of a mix of experienced Teradata users and beginners.
Objectives	<p>This course is designed to get students Teradata utilities ready. Students will understand the Teradata utilities such as BTEQ, FastLoad, MultiLoad, TPump, and FastExport.</p> <p>Students will learn to how use each utility effectively and which utility should be used in which circumstances. We also show them a lot of real-world examples.</p> <p>Topics such as comma delimited files, multiple input files and data conversion works are all covered.</p>

Teradata Utilities – Content

Module 1 - BTEQ

- An Introduction to BTEQ
- Logging on to BTEQ
- Using BTEQ to submit queries
- Interactive Mode and Batch Mode
- BTEQ Conditional Logic
- Using BTEQ to Export Data
- BTEQ EXPORT Example Using Record (DATA) Mode
- BTEQ EXPORT Example Using Field (Report) Mode
- BTEQ IMPORT
- BTEQ Return Codes
- BTEQ Commands
- Session Control Commands
- File Control Commands
- Sequence Control Commands

Module 2 - FastLoad

An Introduction to FastLoad
 How FastLoad Works
 FastLoad Limits
 FastLoad Phases – Acquisition and Application
 FastLoad Commands
 Converting Data Types with FastLoad
 FastLoad RESTART
 FastLoad CHECKPOINT
 Using INMODs with FastLoad

Module 3 - MultiLoad

An Introduction to MultiLoad
 Two MultiLoad Modes: IMPORT and DELETE
 MultiLoad Limits
 Error Tables, Work Tables and Log Tables
 Supported Input Formats
 Phase 1: Preliminary Phase
 Phase 2: DML Transaction Phase
 Phase 3: Acquisition Phase
 Phase 4: Application Phase
 Phase 5: Clean Up Phase
 MultiLoad Commands
 Parameters for .BEGIN IMPORT MLOAD
 Parameters for .BEGIN DELETE MLOAD
 Error Treatment Options for the .DML LABEL Command
 Redefining the INPUT
 What Happens when MultiLoad Finishes
 MultiLoad Statistics
 Troubleshooting MultiLoad Errors
 RESTARTing MultiLoad
 RELEASE MLOAD: When Not to Restart MultiLoad
 MultiLoad and INMODs
 How MultiLoad Compares with FastLoad

Module 4 - TPump

- An Introduction to TPump
- TPump Abilities and Limits
- Supported Input Formats
- TPump Commands and Parameters
- LOAD Parameters IN COMMON with MultiLoad
- .BEGIN LOAD Parameters UNIQUE to TPump
- TPump Output Statistics
- Monitoring TPump
- TPump Error Table
- Common Error Codes
- RESTARTing TPump
- TPump and MultiLoad Comparison

Module 5 – FastExport

- An Introduction to FastExport
- How FastExport Works
- FastExport Fundamentals
- FastExport Supported Operating Systems
- FastExport Support and Task Commands
- FastExport Supported SQL Commands
- FastExport Modes and Formats

Teradata Physical Implementation and Database Design - Overview

Prerequisites	There is no prerequisites for this course.
Duration	This is a 4-day course designed to be highly interactive with the audience.
Courseware	“Teradata Database Design: Giving Detailed Data Flight”.
Audience	The audience will consist of a mix of intermediate and experienced Teradata users.
Objectives	<p>This course details all major aspects of the physical implementation of Teradata and shows users how to tune Teradata for maximum performance.</p> <p>Subjects such as data placement, primary indexes, secondary indexes, joins, how Teradata reads and writes data, manages space, and how performance can be maximized are all covered.</p>

Teradata Physical Implementation and Database Design - Content

Module 1 — Teradata: Brilliant by Design

- The Teradata Parallel Architecture
- A Logical View of the Teradata Architecture
- Components of a Teradata System
- Teradata Linear Scalability
- Parsing Engine (PE)
- Access Module Processor (AMP)
- The BYNET
- Teradata’s Building Block Approach
- Teradata Performance – The Divide and Conquer Theory
- Teradata Data Placement
- PERM and SPOOL Space
- Teradata Data Loads
- Why Teradata is the Database of Choice for Data Warehousing

Module 2 — Data Layout

- Data Distribution Explained
- The Primary Index
- Unique Primary Index (UPI)
- Non-Unique Primary Index
- How Teradata Turns 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
- How Teradata Retrieves Rows
- Row Distribution
- Data Layout Summary

Module 3 — Secondary Indexes

- Secondary Indexes
- Unique Secondary Index (USI)
- Value-Ordered NUSI
- Collect Statistics
- NUSI Bitmapping
- Prototyping indexes with EXPLAIN
- Primary v Secondary Access

Module 4 — Join Strategies

- Join Types vs. Join Strategies
- Merge Join Strategies
- Nested Join
- Hash Join
- Exclusion Join
- Product Joins
- Cartesian Product Join

Module 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
- Column ACCESS in the WHERE Clause
- Data Demographics
- Distinct Values
- Maximum Rows Per Value
- Typical Rows Per Value
- Maximum Rows NULL
- Change Rating
- Extended Logical Data Model Template

Module 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
 Join Access Frequency
 Value Access Frequency
 Step 5 — Pick Secondary Indexes
 USI to Enforce Uniqueness
 USI for High Value Access Frequency
 USI to eliminate Duplicate Row Checking
 Multi-Column NUSI Columns
 Multi-Column NUSI Columns used as a Covered Query
 Value-Ordered NUSIs
 A Strongly Vs Weakly Selective NUSI
 Typical Row Size
 Typical Block Size

Module 7 — Denormalisation

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

Module 8 — Join and Hash Indexes

Join Index Fundamentals
 When to use a Join Index
 Join Indexes versus Views
 Join Indexes versus Summary Tables
 Join Indexes versus Temporary Tables
 Types of Join Indexes
 Single-Table Join Indexes
 Multi-table Join Indexes
 Aggregate Join Index
 Hash Indexes
 Hash Indexes vs. Secondary Indexes
 Hash Indexes vs. Single-Table Join Indexes
 Hash and Join Indexes vs. Base Tables

Module 9 — Explains

How the Teradata Optimiser Explains in Detail
Row Estimate Confidence Levels
Explain Terminology

Teradata Database Administration - Overview

Prerequisites	There are no prerequisites for this course.
Duration	This is a 4-day course designed to be highly interactive with the audience.
Courseware	Teradata Administration.
Audience	The audience will consist of a mix of experienced Teradata users and beginners.
Objectives	This course is taught with a very high hands-on content. We make sure that students understand and are fully prepared to confidently administer the Teradata data warehouse.

Teradata Database Administration - Content

Module 1 - Data Dictionary

Data Dictionary Introduction
 Fallback Protected Tables
 Non-Hashed Tables (AMP-Local)
 Updating Data Dictionary Tables
 System Views
 Supplied Dictionary Views
 Restricted Views
 Selecting Information about Created Objects
 Children View
 Databases View
 Users View
 Tables View
 Tables2 View
 ShowTblChecks View
 ShowColChecks View
 Columns View
 Indices View
 AllTempTable View
 Triggers View

Module 2 - Users, Accounts, Accounting and WINDDI

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

Module 3 - Access Rights

Automatic, Implicit, and Explicit Rights
Tools for finding Access Rights
The GRANT Statement
The REVOKE Statement
Revoking Non-existent Rights
Removing a Level in the Hierarchy
Inheriting Access Rights
The GIVE Statement and Access Rights
System Views for Access Rights
Access Rights Abbreviations
AllRights and UserRights Views
DBC.UserGrantedRights Views

Module 4 - System Access Controls

System Access Control Levels
Teradata Password Encryption
Password Security Features
TDP Exits
Host Logon Processing
GRANT/REVOKE LOGON Statement
DBC.SessionInfo View
DBC.LogOnOff View
DBC.LogonRules View
Access Control Mechanisms
Creating a Simple View
CREATE and EXECUTE MACRO
Using Macros to Insert Data

Module 5 - Access Modelling

BEGIN/END LOGGING Statement
Setting Up Access Logging
Data Access Information Views
AccLogRules Views

Module 6 – Performance Monitoring Tools

Resource Usage Data
Resource Usage Logging
Setting Resource Logging
Collection Costs
Resource Usage Tables
Resource Usage Macros
ResNode Macros
Performance Monitor
Teradata Manager
Dynamic Utilization Charting
PMON
RSSMon

Module 7 - Teradata Utilities

Starting the Database Window (DBW)
Running the DBW
Areas of the DBW
Qryconfig
Qrysessn
Rcvmanager
Showlocks
Showspace.Teradata

Module 8 – Modelling

Logical Data Modeling (Primary Keys)
Logical Data Modeling (Foreign Keys)
Logical Data Model (Normalization)
Extended Logical Model (ELDM)
Physical Data Model
Testing and Production
Capacity Planning
Application Development Cycle Review

Module 9 – Indexing

Primary Indexes
Hash Map
How the Hash Map Distributes Data
How the Hash is used to Locate Data Rows
Row Distribution
Primary Index Considerations
Primary Index Space Utilization
Secondary Indexes
Choosing a Secondary Index or Value Ordered NUSI

Module 10 - Joins and Strategies

Join Types Vs Join Strategies
Simple Join
Merge Join 2 (Primary Index = Column)
Merge Join 3 (Column = Column).
Merge Join 4 (Big Table Small Table Join)
Left Outer Join
Right Outer Join
Full Outer Join
Nested Join
Row Hash Match Scan Join
Exclusion Join
Product Joins
Cartesian Product Join
Using Union to improve Query Performance

Module 11 – Locks

- Locking Modes
- The Lock Manager
- Conflicting Locks
- Locking Modifier
- The NOWAIT Option
- Distinct Pseudo Tables
- Locking Inside Views
- Conflicting Locks with Multi-statement Requests

Module 12 - Data Archiving

- Archive and Recovery Phases
- Recovery Vs FastLoad
- Session Control
- Multiple Sessions
- ARC Statement
- ANALYZE Statement
- ARCHIVE Statement
- Archive Objects
- Archive Levels
- Archive Options
- Indexes Option
- Group Read Lock Option
- Types of Archives
- Database DBC Archive

Module 13 - MultiLoad Statistics

- MultiLoad Statistic Table Create Statement and Column Parameters
- MultiLoad Create Statement with Integrated Statistics Output

Module 14 - Join and Hash Indexes

- Join Index Create Statement
- Hash Index Applications and Configuration
- Hash Index Create Statement

Teradata Application Development – Overview

Prerequisites	There are no prerequisites for this course.
Duration	This is a 3-day course designed to be highly interactive with the audience.
Courseware	Teradata Application Development.
Audience	The audience will consist of a mix of experienced Teradata users and beginners.
Objectives	This course is for application developers who would like to be able to fully understand the application development cycle. Students come away with a well rounded Teradata application development education that will enhance their ability, and their confidence, to perform.

Teradata Application Development - Content

Module 1 - Teradata Architecture

Data Distribution, Hashing, and Index Access
 Teradata Database
 Data Distribution
 Teradata Scalability
 Role of the PE and AMP
 Logical View of Teradata
 Primary Indexes
 How the Hash Distributes Data
 How the Hash is used to Locate Data Rows
 Collect Statistics
 Reconfiguration
 Row ID
 Locating a Row
 Table ID
 Master Index
 Cylinder Index

Module 2 - Temporary Tables

Temp Space
Temporary Tables
Temporary Table Choices
Derived Tables
VOLATILE Temporary Table Example
GLOBAL Temporary Tables

Module 3 - Teradata Parser

Parsing Engine Software Cache
SQL Parser and the Six Main Steps
Request-To-Steps (RTS) Cache
Parser Flow

Module 4 – Modelling

Rows and Columns
Logical Data Modeling (Primary Keys)
Logical Data Modeling (Foreign Keys)
Logical Data Model (Normalization)
Extended Logical Model (ELDM)
Physical Data Model
Testing and Production
Capacity Planning
Application Development Cycle Review

Module 5 – Indexing

Primary Indexes
Hash Map
How the Hash Map Distributes Data
How the Hash is used to Locate Data Rows
Row Distribution
Primary Index Considerations
Primary Index Space Utilization
Secondary Indexes
Choosing a Secondary Index or Value Ordered NUSI

Module 6 – Joins

- Join Types Vs Join Strategies
- Simple Join
- How joins work in Teradata
- Merge Join 2 (Primary Index = Column)
- Merge Join 3 (Column = Column)
- Merge Join 4 (Big Table Small Table Join)
- Left Outer Join
- Right Outer Join
- Full Outer Join
- Nested Join
- Row Hash Match Scan Join
- Exclusion Join
- Product Joins
- Cartesian Product Join
- Using Union to improve Query Performance
- Self Join

Module 7 - Tips And Techniques

- What is a View?
- What is a Macro?
- Collecting Statistics
- Collect Statistics Recommendations
- INSERT SELECT
- Optimized INSERT/SELECT
- Aggregations
- Derived Data

Module 8 – Locks

- Locking Modes
- The Lock Manager
- Conflicting Locks
- Locking Modifier
- The NOWAIT Option
- Distinct Pseudo Tables
- Locking Inside Views
- Conflicting Locks with Multi-statement Requests

Module 9 – OLAP And Hash Functions

- Hashing Functions and HASHROW
- HASHBUCKET and HASHAMP
- HASHBAKAMP

Module 10 – Explain

Explain Overview
Explain Information known to the Optimiser
Optimiser Facts
Row Estimate Confidence Levels
Explain Terminology

Module 11 - Design Considerations For Performance

Design Consideration for Performance
Design Considerations when Planning for Growth
Design Considerations and Optimizer Access Scenarios
Design Considerations and Optimizer Statistics
Design Considerations and Collecting Statistics
Design Considerations and Collected Demographics

Module 12 - Application Development Tools

PP2 The Teradata Preprocessor
PP2 Functionality
Multiple Row Handling
Dynamic SQL
Teradata's Call Level Interface (CLI)
Teradata and ODBC
ODBC Data Sources
Teradata and JDBC
Tips, Tricks, Traps, Techniques

Further Information

Courses are normally run on client premises.

Clients are expected to provide the following:

- Teradata server on which we can load the training database.
- PC workstation for each student (max 2 students per PC).
- QueryMan and ODBC software installed, configured and tested.
- Network connectivity from student workstations to Teradata server.
- DBA resource for training database setup.
- PC compatible projector device.

We will supply the following:

- Training database setup script and instructions.
- Student reference books which each student keeps.
- Courseware.
- Trainer(s).
- Training certificates.

At the end of each course all students are required to complete a course evaluation. We are happy to make copies of each evaluation available to the client.

Please note that all courses are taught in English only. The instructor(s) for each course will be confirmed once dates are agreed.