

Course duration

- 2 days

Course Benefits

- Gain the knowledge to be able to make strategic decisions regarding their Teradata environment.

Course Outline

1. The Teradata Architecture
 1. What is Parallel Processing?
 2. The Basics of a Single Computer
 3. Teradata Parallel Processes Data
 4. Parallel Architecture
 5. The Teradata Architecture
 6. All Teradata Tables are spread across ALL AMPS
 7. Teradata Systems can Add AMPs for Linear Scalability
 8. Understand that Teradata can scale to incredible size
 9. AMPs and Parsing Engines (PEs) live inside SMP Nodes
 10. Each Node is attached via a Network to a Disk Farm
 11. Two SMP Nodes Connected Become One MPP System
 12. There are Many Nodes in a Teradata Cabinet
 13. Inside a Teradata Node
 14. The Boardless BYNET and the Physical BYNET
 15. The Parsing Engine
 16. The AMPs Responsibilities
 17. This is the Visual You Want to Know in order to Understand Teradata
2. The Primary Index
 1. The Primary Index is defined when the table is CREATED
 2. A Unique Primary Index (UPI)
 3. Primary Index in the WHERE Clause - Single-AMP Retrieve
 4. Using EXPLAIN
 5. A Non-Unique Primary Index (NUPI)
 6. Primary Index in the WHERE Clause - Single-AMP Retrieve
 7. Using EXPLAIN in a NUPI Query
 8. A conceptual example of a Multi-Column Primary Index
 9. Primary Index in the WHERE Clause - Single-AMP Retrieve
 10. A conceptual example of a Table with NO PRIMARY INDEX
 11. A Full Table Scan is likely on a table with NO Primary Index
 12. An EXPLAIN that shows a Full Table Scan

13. Table CREATE Examples with four different Primary Indexes
14. What happens when you forget the Primary Index?
15. Why create a table with No Primary Index (NoPI)?
3. Hashing of the Primary Index
 1. The Hashing Formula Facts
 2. The Hash Map determines which AMP will own the Row
 3. The Hash Map determines which AMP will own the Row
 4. Placing rows on the AMP
 5. Placing rows on the AMP Continued
 6. A Review of the Hashing Process
 7. Non-Unique Primary Indexes have Skewed Data
 8. The Uniqueness Value
 9. The Row Hash and Uniqueness Value make up the Row-ID
 10. A Row-ID Example for a Unique Primary Index
 11. A Row-ID Example for a Non-Unique Primary Index (NUPI)
 12. Two Reasons why each AMP Sorts their rows by the Row-ID
 13. AMPs sort their rows by Row-ID to Group like Data
 14. AMPs sort their rows by Row-ID to do a Binary Search
 15. Table CREATE Examples with four different Primary Indexes
 16. Null Values all Hash to the Same AMP
 17. A Unique Primary Index (UPI) Example
 18. A Non-Unique Primary Index (NUPI) Example
 19. A Multi-Column Primary Index Example
 20. A No Primary Index (NoPI) Example
4. Teradata - The Cold Hard Facts
 1. All Teradata Tables are spread across All AMPs
 2. The Table Header and the Data Rows are Stored Separately
 3. An AMP Stores the Rows of a Table inside a Data Block
 4. To Read a Data Block, an AMP Moves the Block into Memory
 5. Nothing is done on disk and everything is done in Memory
 6. Most Taxing thing for an AMP is Moving Blocks into Memory
 7. A Full Table Scan Means All AMPs must Read All Rows
 8. The "Achilles Heel and slowest process is Block Transfer
 9. Each Table has a Primary Index
 10. A Query Using the Primary Index is a Single AMP Retrieve.
 11. As Rows are added a Data Block will Eventually Split
 12. A Full Table Scan Means All AMPs must Read All Blocks
 13. A Primary Index Query uses a Single AMP and Single Block
 14. Each AMP Can Have Many Blocks for a Single Table
 15. A Full Table Scan Means All AMPs must Read All Blocks
 16. Synchronized Scan (Sync Scan)
 17. EXPLAIN Using a Synchronized Scan
 18. Intelligent Memory (Teradata V14.10)
 19. Teradata V14.10 Intelligent Memory Gives Data a Temperature
 20. Data deemed VeryHot stays in each AMP's Intelligent Memory
 21. Intelligent Memory Stays in Memory
 22. What is the Goal of a Teradata Physical Database Design?

5. Inside the AMPs Disk

1. Rows are Stored in Data Blocks which are stored in Cylinders
2. An AMP's rows are stored inside a Data Block in a Cylinder
3. An AMP's Master Index is used to find the Right Cylinder
4. The Row Reference Array (RRA) Does the Binary Search?
5. A Block Splits into Two Blocks at Maximum Block Size
6. Data Blocks Maximum Block Size has Changed (V14.10)
7. The New Block Split with Teradata V14.10
8. The Block Split with Even More Detail in Teradata V14.10
9. Teradata V14.10 Block Split Defaults
10. There is One Master Index and Thousands of Cylinder Indexes
11. Blocks Continue to Split as Tables Grow Larger
12. FYI – Some Advanced Information about Data Block Headers
13. A top down view of Cylinders
14. There are Hot, Warm, and Cold Cylinders
15. Cylinders are used for Perm, Spool, Temp, and Journals
16. Each AMP has Their Own Master Index
17. Each Cylinder on an AMP has a Cylinder Index
18. A More Detailed Illustration of the Master Index
19. A Real-World View of the Master Index
20. An Even More Realistic View of an AMP's Master Index
21. The Cylinder Index
22. An Even More Realistic View of a Cylinder Index
23. How a Query using the Primary Index works
24. How the AMPs Do a Full Table Scan
25. How an AMP Reads Using a Primary Index

6. Partition Primary Index (PPI) Tables

1. The Concept behind Partitioning a Table
2. Creating a PPI Table with Simple Partitioning
3. A Visual Display of Simple Partitioning
4. An SQL Example that explains Simple Partitioning
5. Creating a PPI Table with RANGE_N Partitioning per Month
6. A Visual of One Year of Data with Range_N per Month
7. An SQL Example explaining Range_N Partitioning per Month
8. A Partition # and Row-ID = Row Key
9. An AMP Stores its Rows Sorted in only Two Different Ways
10. Creating a PPI Table with RANGE_N Partitioning per Day
11. A Visual of Range_N Partitioning Per Day
12. An SQL Example that explains Range_N Partitioning per Day
13. Creating a PPI Table with RANGE_N Partitioning per Week
14. A Visual of Range_N Partitioning Per Week
15. SQL Example that explains Range_N Partitioning per Week
16. A Clever Range_N Option
17. Creating a PPI Table with CASE_N
18. A Visual of Case_N Partitioning
19. An SQL Example that explains CASE_N Partitioning
20. How many partitions do you see?

21. Number of PPI Partitions Allowed
22. How many partitions do you see?
23. NO CASE and UNKNOWN Partitions Together
24. A Visual of Case_N Partitioning
25. Combining Older Data and Newer Data in PPI
26. A Visual for Combining Older Data and Newer Data in PPI
27. The SQL on Combining Older Data and Newer Data in PPI
28. Multi-Level Partitioning Combining Range_N and Case_N
29. A Visual of Multi-Level Partitioning
30. The SQL on a Multi-Level Partitioned Primary Index
31. NON-Unique Primary Indexes (NUPI) in PPI
32. PPI Table with a Unique Primary Index (UPI)
33. Tricks for Non-Unique Primary Indexes (NUPI)
34. Character Based PPI for RANGE_N
35. A Visual for Character-Based PPI for RANGE_N
36. The SQL on Character-Based PPI for RANGE_N
37. Character-Based PPI for CASE_N
38. Dates and Character-Based Multi-Level PPI
39. TIMESTAMP Partitioning
40. Using CURRENT_DATE to define a PPI
41. ALTER to CURRENT_DATE the next year
42. ALTER to CURRENT_DATE with Save
43. Altering a PPI Table to Add or Drop Partitions
44. Deleting a Partition
45. Deleting a Partition and saving its contents
46. Using the PARTITION Keyword in your SQL
47. SQL for RANGE_N
48. SQL for CASE_N

7. Columnar Tables

1. Columnar Tables have NO Primary Index
2. This is NOT a NoPI Table
3. NoPI Tables Spread rows across all-AMPs Evenly
4. NoPI Tables used as Staging Tables for Data Loads
5. NoPI Table Capabilities
6. NoPI Table Restrictions
7. What does a Columnar Table look like?
8. Comparing Normal Table vs. Columnar Tables
9. Columnar Table Fundamentals
10. Example of Columnar CREATE Statement
11. Columnar can move just One Container to Memory
12. Containers on AMPs match up perfectly to rebuild a Row
13. Indexes can be used on Columns (Containers)
14. Indexes can be used on Columns (Containers)
15. Visualize a Columnar Table
16. Single-Column vs. Multi-Column Containers
17. Comparing Normal Table vs. Columnar Tables
18. Columnar Row Hybrid CREATE Statement

19. Columnar Row Hybrid Example
20. Columnar Row Hybrid Query Example
21. Review of Row-Based Partition Primary Index (PPI)
22. Visual of Row Partitioning (PPI Tables) by Month
23. CREATE Statement for both Row and Column Partition
24. Visual of Row Partitioning (PPI Tables) and Columnar
25. How to Load into a Columnar Table
26. Columnar NO AUTO COMPRESS
27. Auto Compress in Columnar Tables
28. Auto Compress Techniques in Columnar Tables
29. When and When NOT to use Columnar Tables
30. Did you know?

8. Space

1. When your System Arrives, there is only User named DBC
2. USER DBC
3. First Assignment is to create another User just under DBC
4. USER DBC
5. Perm and Spool Space
6. Perm Space is for Permanent Tables
7. Spool Space is work space that builds a User's Answer Sets
8. Spool Space is in an AMP's Memory and on its Disk
9. Users are Assigned Spool Space Limits
10. What is the Purpose of Spool Limits?
11. Why did my query Abort and say "Out of Spool"?
12. How can Skewed Data cause me to run "Out of Spool"?
13. Why did my Join cause me to run "Out of Spool"?
14. Finding out how much Space you have
15. Space per AMP on all tables in a Database shows Skew
16. What does my system look like when it first arrives?
17. DBC owns all the PERM Space in the system on day one
18. DBC's First Assignment is Spool Space
19. DBC's 2nd Assignment is to CREATE Users and Databases
20. The Teradata Hierarchy Begins
21. The Teradata Hierarchy Continues
22. Differences between PERM and SPOOL
23. Databases, Users, and Views
24. What are Similarities between a DATABASE and a USER?
25. What is the Difference between a DATABASE and a USER?
26. Objects that take up PERM Space
27. A Series of Quizzes on Adding and Subtracting Space

9. The User Environment

1. DBC is the only user when the system first arrives
2. DBC will Create Databases and Give them Space
3. DBC will create some initial Users
4. A Typical Teradata Environment
5. What are Similarities between a DATABASE and a USER?
6. Roles

7. Create a Role and then Assign that Role Its Access Rights
8. Create a User and Assign them a Default Role
9. Granting Access Rights
10. There are Three Types of Access Rights
11. Description of the Three Types of Access Rights
12. Profiles
13. Creating a Profile and a User
14. ProfileInfoVX, RoleMembers, RoleInfo and UserRoleRights
15. Accounts and their Associated Priorities
16. Creating a User with Multiple Account Priorities
17. Account String Expansion (ASE)
18. The DBC.AMPUsage View
19. Teradata TASM provides a User Traffic System
20. Teradata Viewpoint
10. Secondary Indexes
 1. Creating a Unique Secondary Index (USI)
 2. What is in a Unique Secondary Index (USI) Subtable?
 3. A Unique Secondary Index (USI) Subtable is hashed
 4. How the Parsing Engine uses the USI Subtable
 5. A USI is a Two-AMP Operation
 6. Creating a Non-Unique Secondary Index (NUSI)
 7. What is in a Unique Secondary Index (USI) Subtable?
 8. Non-Unique Secondary Index (NUSI) Subtable is AMP Local
 9. How the Parsing Engine uses the NUSI Subtable
 10. Creating a Value-Ordered NUSI
 11. The Hash Map Determines which AMP will own the Row
 12. A Unique Primary Index Spreads the Data Evenly
 13. A Picture with a Base Table, USI, and NUSI Subtable
 14. A Query Using an USI Only Moves Two Blocks
 15. A Query Using A NUSI Always Uses All AMPs
 16. Two Non-Unique Secondary Indexes (NUSI) on a Table
 17. A NUSI BITMAP Query (1 of 3)
 18. A NUSI BITMAP Theory (2 of 3)
 19. A NUSI Bitmap in Action (3 of 3)
 20. A Brilliant Technique for a Unique Secondary Index
 21. The USI for Partitioned Tables Points to the Row Key
 22. A Brilliant Technique for a Non-Unique Secondary Index
 23. The NUSI for Partitioned Tables Points to the Row Key
 24. How the PE Decides on the NUSI or the Full Table Scan
 25. Multiple Choice DBA
 26. What are the Big Four Tactical Queries?
11. Temporal Tables Create Functions
 1. Three types of Temporal Tables
 2. CREATING a Bi-Temporal Table
 3. PERIOD Data Types
 4. Bi-Temporal Data Type Standards
 5. Bi-Temporal Example – Tera-Tom buys!

6. A Look at the Temporal Results
7. Bi-Temporal Example – Tera-Tom Sells!
8. Bi-Temporal Example – How the data looks!
9. Normal SQL for Bi-Temporal Tables
10. NONSEQUENCED SQL for Temporal Tables
11. AS OF SQL for Temporal Tables
12. NONSEQUENCED for Both
13. Creating Views for Temporal Tables
14. Bi-Temporal Example – Socrates is DELETED!
12. How Joins Work Internally
 1. The Joining of Two Tables
 2. Teradata Moves Joining Rows to the Same AMP
 3. Imagine Joining Two NoPI Tables that have No Primary Index
 4. Both Tables are redistributed to Join Rows on the Same AMP
 5. How do you join if One Table is Big and One Table is Small?
 6. Duplicate the Small Table on Every AMP (like a mirror)
 7. What Could You Do If Two Tables Joined 1000 Times a Day?
 8. Joining Two Tables with the same PK/FK Primary Index
 9. A Join with No Redistribution or Duplication
 10. A Performance Tuning Technique for Large Joins
 11. The Joining of Two Tables with an Additional WHERE Clause
 12. An Example of the Fastest Join Possible
 13. Using a Simple Volatile Table
 14. A Volatile Table with a Primary Index
 15. Using a Simple Global Temporary Table
 16. Two Brilliant Techniques for Global Temporary Tables
 17. The Joining of Two Tables Using a Global Temporary Table
 18. Teradata V14.10 Join Feature PRPD
13. Join Indexes
 1. Creating a Multi-Table Join Index
 2. Visual of a Join Index
 3. Outer Join Multi-Table Join Index
 4. Visual of a Left Outer Join Index
 5. Compressed Multi-Table Join Index
 6. A Visual of a Compressed Multi-Table Join Index
 7. Creating a Single-Table Join Index
 8. Conceptual of a Single Table Join Index on an AMP
 9. Single Table Join Index Great For LIKE Clause
 10. Single Table Join Index with Value Ordered NUSI
 11. Aggregate Join Indexes
 12. Compressed Single-Table Join Index
 13. Aggregate Join Index
 14. New Aggregate Join Index (Teradata V14.10)
 15. Sparse Join Index
 16. A Global Multi-Table Join Index
 17. Creating a Hash Index
 18. Join Index Details

14. Collect Statistics

1. The Teradata Parsing Engine (Optimizer) is Cost Based
2. The Purpose of Collect Statistics
3. When Teradata Collects Statistics it creates a Histogram
4. The Interval of the Collect Statistics Histogram
5. What to COLLECT STATISTICS On?
6. Why Collect Statistics?
7. How do you know if Statistics were collected on a Table?
8. A Huge Hint that No Statistics Have Been Collected
9. The Basic Syntax for COLLECT STATISTICS
10. COLLECT STATISTICS Examples for a better Understanding
11. The New Teradata V14 Way to Collect Statistics
12. Where Does Teradata Keep the Collected Statistics?
13. The Official Syntax for COLLECT STATISTICS
14. How to Recollect STATISTICS on a Table
15. Teradata Always Does a Random AMP Sample
16. Random Sample is kept in the Table Header in FSG Cache
17. Multiple Random AMP Samplings
18. How a Random AMP gets a Table Row count
19. Random AMP Estimates for NUSI Secondary Indexes
20. USI Random AMP Samples are Not Considered
21. There's No Random AMP Estimate for Non-Indexed Columns
22. The PE's Plan if No Statistics Were Collected?
23. Stale Statistics Detection and Extrapolation
24. Extrapolation for Future Dates
25. How to Copy a Table with Data and the Statistics?
26. How to Copy a Table with NO Data and the Statistics?
27. COLLECT STATISTICS Directly From another Table
28. When to COLLECT STATISTICS Using only a SAMPLE
29. Examples of COLLECT STATISTICS Using only a SAMPLE
30. Examples of COLLECT STATISTICS For V14
31. How to Collect Statistics on a PPI Table on the Partition
32. Teradata V12 and V13 Statistics Enhancements
33. Teradata V14 Statistics Enhancements
34. Teradata V14 Summary Statistics
35. Teradata V14 MaxValueLength
36. Teradata V14 MaxIntervals
37. Teradata V14 Sample N Percent
38. Teradata V14.10 Statistics Collection Improvements
39. Teradata V14.10 Statistics Collection Improvements
40. Teradata V14.10 AutoStats feature
41. Teradata Statistics Wizard

15. Temporary Tables

1. There are three types of Temporary Tables
2. CREATING A Derived Table
3. Naming the Derived Table
4. Aliasing the Column Names in the Derived Table

5. Most Derived Tables Are Used To Join To Other Tables
6. Multiple Ways to Alias the Columns in a Derived Table
7. Our Join Example with a Different Column Aliasing Style
8. Column Aliasing Can Default for Normal Columns
9. CREATING A Derived Table using the WITH Command
10. Our Join Example With the WITH Syntax
11. The Same Derived Query shown Three Different Ways
12. Clever Tricks on Aliasing Columns in a Derived Table
13. A Derived Table lives only for the lifetime of a single query
14. An Example of Two Derived Tables in a Single Query
15. WITH RECURSIVE Derived Table
16. Defining the WITH Recursive Derived Table
17. Looping Through the Recursive Derived Table
18. Looping Through a Second Time
19. Looping Through a Third Time
20. Looping Through and Adding Nothing Ends the Loop
21. Looping Through the WITH Recursive Derived Table
22. Creating a Volatile Table
23. You Populate a Volatile Table with an INSERT/SELECT
24. The Three Steps to Use a Volatile Table
25. Why Would You Use the ON COMMIT DELETE ROWS?
26. The HELP Volatile Table Command Shows your Volatiles
27. A Volatile Table with a Primary Index
28. The Joining of Two Tables Using a Volatile Table
29. You Can Collect Statistics on Volatile Tables
30. The New Teradata V14 Way to Collect Statistics
31. Four Examples of Creating a Volatile Table Quickly
32. Four Advanced Examples of Creating a Volatile Table Quickly
33. Creating Partitioned Primary Index (PPI) Volatile Tables
34. Using a Volatile Table to Get Rid of Duplicate Rows
35. Using a Simple Global Temporary Table
36. Two Brilliant Techniques for Global Temporary Tables
37. The Joining of Two Tables Using a Global Temporary Table
38. CREATING A Global Temporary Table
16. Teradata Load Utilities Introduction
 1. The Teradata Utilities
 2. Block Level Utilities
 3. Row Level Utilities
 4. Fast Path Inserts Using Insert/Select
 5. Fast Path Deletes
 6. Freespace Percent
 7. Referential Integrity and Load Utility Solutions
 8. Teradata has a No Primary Index Table called a NoPI Table
 9. This is NOT Necessarily a NoPI Table
 10. NoPI Tables Spread rows across all-AMPs Evenly
 11. NoPI Tables used as Staging Tables for Data Loads
 12. NoPI Table Capabilities

13. NoPI Table Restrictions
 14. Why Would a NoPI Table have a Row-ID?
 15. BTEQ – Batch Teradata Query Tool
 16. How to Logon to BTEQ in Interactive Mode
 17. Running Queries in BTEQ in Interactive Mode
 18. BTEQ Commands vs BTEQ SQL Statements
 19. WITH BY Command for Subtotals
 20. WITH Command for a Grand Total
 21. WITH and WITH BY Together for Subtotals and Grand Totals
 22. How to Logon to BTEQ in a SCRIPT
 23. Running Queries in BTEQ through a Batch Script
 24. Running a BTEQ Batch Script through the Command Prompt
 25. Running a BTEQ Batch Script through the Run Command
 26. Using Nexus to Build Your BTEQ Scripts
 27. Using Nexus to Build Your BTEQ Scripts
 28. FastLoad
 29. Block Level Utility Limits
 30. FastLoad has Two Phases
 31. FastLoad Phase 1
 32. FastLoad Phase 2
 33. A Sample FastLoad Script Created by Nexus SmartScript
 34. Executing the FastLoad Script
 35. The Nexus SmartScript Easily Builds Your Utilities
 36. The Nexus SmartScript FastLoad Builder
 37. Create and Execute Your FastLoad Scripts with Nexus
 38. MultiLoad
 39. Block Level Utility Limits
 40. MultiLoad has Five Phases
 41. MultiLoad has IMPORT and DELETE Tasks
 42. A Sample MultiLoad Script Created by Nexus SmartScript
 43. TPump
 44. TPump is NOT a Block Level Utility and has No Limits
 45. Limitations of TPump
 46. A Sample TPump Script Created by Nexus SmartScript
 47. FastExport
 48. New Rules for Block Utilities
 49. A Sample FastExport Script Created by Nexus SmartScript
 50. FastExport by Default places Null Indicators in Output
 51. A Sample FastExport Script Created by Nexus SmartScript
 52. What is TPT?
 53. TPT Producers Create Streams and Consumers Write Them
 54. The Four Major Operators of TPT
 55. TPT can read from multiple source files in Parallel
 56. TPT can have more Operators than Consumers
 57. TPT Operators and their Equivalent Load Utility
 58. How to Run a TPT Script
17. Top SQL Commands Cheat Sheet

1. SELECT All Columns from a Table and Sort
2. Select Specific Columns and Limiting the Rows
3. Changing your Default Database
4. Keywords that describe you
5. Select TOP Rows in a Rank Order
6. A Sample number of rows
7. Getting a Sample Percentage of rows
8. Find Information about a Database
9. Find information about a Table
10. Using Aggregates
11. Performing a Join
12. Performing a Join using ANSI Syntax
13. Using Date, Time and Timestamp
14. Using Date Functions
15. Using the System Calendar
16. Using the System Calendar in a Query
17. Formatting Data
18. Using Rank
19. Using a Derived Table
20. Using a Subquery
21. Correlated Subquery
22. Using Substring
23. Basic CASE Statement
24. Advanced CASE Statement
25. Using an Access Lock in your SQL
26. Collect Statistics
27. CREATING a Volatile Table with a Primary Index
28. CREATING a Volatile Table that is Partitioned (PPI)
29. CREATING a Volatile Table that is deleted after the Query
30. Finding the Typical Rows per Value for specific column
31. Finding out how much Space you have
32. How much Space you have Per AMP
33. Finding your Space
34. Finding Space Skew in Tables in a Database
35. Finding the Number of rows per AMP for a Column
36. Finding Account Information
37. Ordered Analytics

Class Materials

Each student will receive a comprehensive set of materials, including course notes and all the class examples.