

Teradata Custom for Amdocs

Course Summary

Description

This course is designed as a 3-day Teradata course to quickly and efficiently bring Teradata students to an entirely new level. Students will start on day one with an intense understanding of the entire Teradata architecture. Day two will combine the architecture with the advanced portions of the SQL with performance tuning in mind to enhance even the best SQL users. Day three will strongly examine the TASM and Viewpoint features so students can understand system administration and monitoring of queries at the deepest level. The three most important aspects of Teradata are the architecture, SQL capabilities and the system administration. Tera-Tom will be teaching this course personally so there will be an enormous amount of hands-on work as well as an enormous opportunity for those new to Teradata and those who already have experience.

Topics

- The Teradata Architecture
- Primary Index
- Hashing of the Primary Index
- Space
- Partition Primary Index (PPI) Tables Subquery
- Secondary Indexes
- Columnar Tables
- Temporal Tables Create Functions
- How Joins work internally
- The TOP Command
- Aggregation Function
- Temporary Tables
- OLAP Functions
- Date Functions
- Join Functions

Duration

Three days

Teradata Custom for Amdocs

Course Outline

Day 1 – Teradata Architecture

I. The Teradata Architecture

- A. Parallel Processing
- B. The Teradata Architecture
- C. All Teradata Tables are spread across All AMPs
- D. Each Table has a Column that is the Primary Index
- E. A Full Table Scan uses All-AMPs in Parallel
- F. Knowing the Primary Index of a Table is Vital
- G. Teradata Systems can Add AMPs for Linear Scalability
- H. Using quantifiers versus IN

II. Primary Index

- A. The Primary Index is defined when the table is CREATED
- B. A Unique Primary Index (UPI)
- C. Primary Index in the WHERE Clause - Single-AMP Retrieve
- D. A Non-Unique Primary Index (NUPI)
- E. Primary Index in the WHERE Clause - Single-AMP Retrieve
- F. A conceptual example of a Multi-Column Primary Index
- G. A Full Table Scan is likely on a table with NO Primary Index
- H. Table CREATE Examples with four different Primary Indexes
- I. What happens when you forget the Primary Index?
- J. Why create a table with No Primary Index (NoPI)?

III. Hashing of the Primary Index

- A. The Hashing Formula Facts
- B. The Hash Map Determines which AMP will own the Row
- C. Placing rows on the AMP
- D. A Review of the Hashing Process
- E. Non-Unique Primary Indexes have Skewed Data
- F. The Uniqueness Value
- G. The Row Hash and Uniqueness Value make up the Row-ID
- H. A Row-ID Example for a Unique Primary Index
- I. A Row-ID Example for a Non-Unique Primary Index(NUPI)
- J. Two Reasons why each AMP Sorts their rows by the Row-ID

- K. AMPs sort their rows by Row-ID to Group Like Data
- L. AMPs sort their rows by Row-ID to do a Binary Search
- M. Table CREATE Examples with four different Primary Indexes
- N. Null Values all Hash to the Same AMP
- O. A Unique Primary Index (UPI) Example
- P. A Non-Unique Primary Index (NUPI) Example
- Q. A Multi-Column Primary Index Example
- R. A No Primary Index (NoPI) Example

IV. Space

- A. Perm and Spool Space
- B. Perm Space is for Permanent Tables
- C. Spool Space is work space that builds a User's Answer Sets
- D. Spool Space is in an AMPs memory and on its Disk
- E. USERS are Assigned Spool Space Limits
- F. What is the Purpose of Spool Limits?
- G. Why did my query Abort and say "Out of Spool"?
- H. How can Skewed Data cause me to run "Out of Spool"?
- I. How come my Join caused me to run "Out of Spool"?
- J. What does my system look like when it first arrives?
- K. DBC owns all the PERM Space in the system on day one
- L. DBC's First Assignment is Spool Space
- M. DBC's 2nd Assignment is to CREATE Users and Databases
- N. The Teradata Hierarchy Begins
- O. Differences between PERM and SPOOL
- P. Databases, Users, and Views
- Q. What are Similarities between a DATABASE and a USER?
- R. What is the Difference between a DATABASE and a USER?
- S. Objects that take up PERM Space
- T. A Series of Quizzes on Adding and Subtracting Space
- U. Space Transfer Quiz
- V. Drop Space Quiz

**Teradata
Custom for Amdocs**

Course Outline (cont'd)

V. Partition Primary Index (PPI) Tables Subquery

- A. The Concept behind Partitioning a Table
- B. Review of A Unique Primary Index (UPI)
- C. Primary Index in the WHERE Clause - Single-AMP Retrieve
- D. Review of A Non-Unique Primary Index (NUPI)
- E. Primary Index in the WHERE Clause - Single-AMP Retrieve
- F. Review of a Multi-Column Primary Index
- G. Primary Index in the WHERE Clause - Single-AMP Retrieve
- H. Creating a PPI Table with Simple Partitioning
- I. A Visual Display of Simple Partitioning
- J. An SQL Example that explains Simple Partitioning
- K. Creating a PPI Table with RANGE_N Partitioning per Day
- L. A Visual of Range_N Partitioning Per Day
- M. An SQL Example that explains Range_N Partitioning per Day
- N. Creating a PPI Table with RANGE_N Partitioning per Week
- O. A Visual of Range_N Partitioning Per Week
- P. An SQL Example that explains Range_N Partitioning per Day
- Q. Creating a PPI Table with RANGE_N Partitioning per Month
- R. A Visual of One Year of Data with Range_N Per Month
- S. An SQL Example explaining Range_N Partitioning per Month
- T. Creating a PPI Table with CASE_N
- U. A Visual of Case_N Partitioning
- V. An SQL Example that explains Range_N Partitioning per Day
- W. How many partitions do you see?
- X. Number of PPI Partitions Allowed
- Y. How many partitions do you see?
- Z. NO CASE and UNKNOWN Partitions Together
- AA. A Visual of Case_N Partitioning
- BB. Combining Older Data and Newer Data in PPI
- CC. A Visual for Combining Older Data and Newer Data in PPI
- DD. The SQL on Combining Older Data and Newer Data in PPI
- EE. Multi-Level Partitioning Combining Range_N and Case_N
- FF. A Visual of Multi-Level Partitioning
- GG. The SQL on a Multi-Level Partitioned Primary Index

- HH. NON-Unique Primary Indexes (NUPI) in PPI
- II. PPI Table with a Unique Primary Index (UPI)
- JJ. Tricks for Non-Unique Primary Indexes (NUPI)
- KK. Character Based PPI for RANGE_N
- LL. A Visual for Character Based PPI for RANGE_N
- MM. The SQL on Character Based PPI for RANGE_N
- NN. Character Based PPI for CASE_N
- OO. Dates and Character Based Multi-Level PPI
- PP. TIMESTAMP Partitioning
- QQ. Using CURRENT_DATE to define a PPI
- RR. ALTER to CURRENT_DATE the next year
- SS. ALTER to CURRENT_DATE with Save
- TT. Altering a PPI Table to Add or Drop Partitions
- UU. Deleting a Partition
- VV. Deleting a Partition and Saving its contents
- WW. Using the PARTITION Keyword in your SQL
- XX. SQL for RANGE_N
- YY. SQL for CASE_N

VI. Secondary Indexes

- A. Review of a Unique Primary Index (UPI)
- B. Primary Index in the WHERE Clause - Single-AMP Retrieve
- C. Review of a Non-Unique Primary Index (NUPI)
- D. Primary Index in the WHERE Clause - Single-AMP Retrieve
- E. Creating a Unique Secondary Index (USI)
- F. What is in a Unique Secondary Index (USI) Subtable?
- G. A Unique Secondary Index (USI) Subtable is Hashed
- H. How the Parsing Engine uses the USI Subtable?
- I. An USI is a Two-AMP Operation
- J. Review of a Non-Unique Primary Index (NUPI)
- K. Creating a Non-Unique Secondary Index (NUSI)
- L. What is in a Unique Secondary Index (USI) Subtable?
- M. Non-Unique Secondary Index (NUSI) Subtable is AMP Local
- N. How the Parsing Engine uses the NUSI Subtable?

VII. Columnar Tables

- A. Columnar Tables have NO Primary Index
- B. This is NOT a NoPI Table
- C. NoPI Tables Spread rows across all-AMPs Evenly
- D. NoPI Tables used as Staging Tables for Data Loads

**Teradata
Custom for Amdocs**

Course Outline (cont'd)

- E. NoPI Table Capabilities
 - F. NoPI Table Restrictions
 - G. What does a Columnar Table look like?
 - H. Comparing Normal Table Vs Columnar Tables
 - I. Columnar Table Fundamentals
 - J. Example of Columnar CREATE Statement
 - K. Columnar can move just One Container to Memory
 - L. Containers on AMPs match up Perfectly to rebuild a Row
 - M. Indexes can be used on Columns (Containers)
 - N. Visualize a Columnar Table
 - O. Single-Column Vs. Multi-Column Containers
 - P. Comparing Normal Table Vs. Columnar Tables
 - Q. Columnar Row Hybrid CREATE Statement
 - R. Columnar Row Hybrid Example
 - S. Review of Row Based Partition Primary Index (PPI)
 - T. Visual of Row Partitioning (PPI Tables) by Month
 - U. CREATE Statement for both Row and Column Partition
 - V. Visual of Row Partitioning (PPI Tables)and Columnar
 - W. How to Load into a Columnar Table
 - X. Columnar NO AUTO COMPRESS
 - Y. Auto Compress in Columnar Tables
 - Z. Auto Compress Techniques in Columnar Tables
 - AA. When and When NOT to use Columnar Tables
- C. If the Join Condition is the Primary Index no Movement
 - D. How the Parsing Engine Decides on a Join Plan
 - E. Quiz – Redistribute the Employees by their Dept_No
 - F. Quiz – Employees Dept_No landed on AMP with Matches
 - G. When Rows are on the same AMP they can be Joined
 - H. Redistribution and then a Row Hash Match Scan
 - I. Quiz – Redistribute the Orders to the Proper AMP
 - J. Answer to Redistribute the Employees by their Dept_No Quiz
 - K. A Visual of the Join in Action
 - L. The Big Table Small Table Join causes Duplication
 - M. Visual of Duplication of the Smaller Table across All-AMPs
 - N. Duplication of the Smaller Table across All-AMPs
 - O. A Visual of Duplication of the Smaller Table on a Single AMP

Day 2 – Advanced SQL

VIII. Temporal Tables Create Functions

- A. Three types of Temporal Tables
- B. CREATING a Bi-Temporal Table
- C. PERIOD Data Types
- D. Bi-Temporal Data Type Standards
- E. Bi-Temporal Example – Tera-Tom buys!
- F. A Look at the Temporal Results
- G. Bi-Temporal Example – Tera-Tom Sells!
- H. Bi-Temporal Example – How the data looks!
- I. Normal SQL for Bi-Temporal Tables
- J. NONSEQUENCED SQL for Temporal Tables
- K. AS OF SQL for Temporal Tables
- L. NONSEQUENCED for Both
- M. NONSEQUENCED for Both
- N. Property_Owners Before DELETE on April 1st

IX. How Joins work internally

- A. Teradata Join Quiz
- B. Teradata Join Quiz Answer

X. The TOP Command

- A. TOP Command
- B. TOP Command is brilliant when ORDER BY Used
- C. The TOP Command WITH TIES
- D. How the TOP Command WITH TIES Decides
- E. The TOP Command will NOT work with Certain Commands

XI. Aggregation Function

- A. Quiz – You calculate the Answer Set in your own Mind
- B. Answer – You calculate the Answer Set in your own Mind
- C. The 3 Rules of Aggregation
- D. There are Five Aggregates
- E. Quiz – How many rows come back?
- F. Troubleshooting Aggregates
- G. GROUP BY when Aggregates and Normal Columns Mix
- H. GROUP BY Delivers one row per Group
- I. GROUP BY Dept_No or GROUP BY 1 the same thing
- J. Limiting Rows and Improving Performance with WHERE

**Teradata
Custom for Amdocs
Course Outline (cont'd)**

- K. WHERE Clause in Aggregation limits unneeded Calculations
- L. Keyword HAVING tests Aggregates after they are Totaled
- M. Keyword HAVING is like an Extra WHERE Clause for Totals
- N. Three types of Advanced Grouping
- O. GROUP BY Grouping Sets
- P. GROUP BY Rollup
- Q. GROUP BY Rollup Result Set
- R. GROUP BY Cube
- S. GROUP BY CUBE Result Set
- T. OUTER JOIN with Additional AND Clause
- U. OUTER JOIN with Additional WHERE Clause
- V. Testing Your Knowledge
- W. Final Answer to Test Your Knowledge on Aggregates

XII. Temporary Tables

- A. There are Three types of Temporary Tables
- B. CREATING A Derived Table
- C. Naming the Derived Table
- D. Aliasing the Columns of the Derived Table
- E. Multiple Ways to Alias the Columns in a Derived Table
- F. CREATING A Derived Table using the WITH Command
- G. Naming the Derived Table using the WITH Command
- H. Naming the Derived Table Columns using WITH
- I. The Same Derived Query shown Three Different Ways
- J. A Derived Table that Joins to an Existing Table
- K. Quiz - Answer the Questions
- L. Clever Tricks on Aliasing Columns in a Derived Table
- M. A Derived Table lives only for the lifetime of a single query
- N. An Example of Two Derived Tables in a Single Query
- O. Creating a Volatile Table
- P. You Populate a Volatile Table with an INSERT/SELECT
- Q. The Three Steps to Use a Volatile Table
- R. The HELP Volatile Table Command Shows your Volatiles
- S. CREATING A Global Temporary Table
- T. Populating A Global Temporary Table with INSERT/SELECT

XIII. OLAP Functions

- A. On-Line Analytical Processing (OLAP) or Ordered Analytics
- B. Cumulative Sum (CSUM) Command and how OLAP Works
- C. OLAP Commands always Sort (ORDER BY) in the Command
- D. Calculate the Cumulative Sum (CSUM) after Sorting the Data
- E. The OLAP Major Sort Key
- F. The OLAP Major Sort Key and the Minor Sort Key(s)
- G. Troubleshooting OLAP – My Data isn't coming back Correct
- H. GROUP BY in Teradata OLAP Syntax Resets on the Group
- I. CSUM the Number 1 to get a Sequential Number
- J. A Single GROUP BY Resets each OLAP with Teradata Syntax
- K. A Better Choice – The ANSI Version of CSUM
- L. The ANSI Version of CSUM – The Sort Explained
- M. The ANSI CSUM – Rows Unbounded Preceding Explained
- N. The ANSI CSUM – Making Sense of the Data
- O. The ANSI CSUM – Making Even More Sense of the Data
- P. The ANSI CSUM – The Major and Minor Sort Key(s)
- Q. The ANSI CSUM – Getting a Sequential Number
- R. Troubleshooting The ANSI OLAP on a GROUP BY
- S. The ANSI OLAP – Reset with a PARTITION BY Statement
- T. PARTITION BY only Resets a Single OLAP not ALL of them
- U. The Moving Average (MAVG) and Moving Window
- V. How the Moving Average is Calculated
- W. How the Sort works for Moving Average (MAVG)
- X. GROUP BY in the Moving Average does a Reset
- Y. Quiz – Can you make the Advanced Calculation in your mind?
- Z. Answer to Quiz for the Advanced Calculation in your mind?
- AA. Quiz – Write that Teradata Moving Average in ANSI Syntax
- BB. Both the Teradata Moving Average and ANSI Version

**Teradata
Custom for Amdocs**

Course Outline (cont'd)

- CC. The ANSI Moving Window is Current Row and Preceding
- DD. How ANSI Moving Average Handles the Sort
- EE. Quiz – How is that Total Calculated
- FF. Answer to Quiz – How is that Total Calculated?
- GG. Quiz – How is that 4th Row Calculated?
- HH. Answer to Quiz – How is that 4th Row Calculated?
- II. Moving Average every 3-rows Vs. a Continuous Average
- JJ. Partition By Resets an ANSI OLAP
- KK. The Moving Difference (MDIFF)
- LL. Moving Difference (MDIFF) Visual
- MM. Moving Difference using ANSI Syntax
- NN. Moving Difference using ANSI Syntax with Partition By
- OO. Trouble Shooting the Moving Difference (MDIFF)
- PP. The RANK Command
- QQ. How to get Rank to Sort in Ascending Order
- RR. Two ways to get Rank to Sort in Ascending Order
- SS. RANK using ANSI Syntax Defaults to Ascending Order
- TT. Getting RANK using ANSI Syntax to Sort in DESC Order
- UU. RANK() OVER and PARTITION BY with a QUALIFY
- VV. QUALIFY and WHERE
- WW. Quiz – How can you simplify the QUALIFY Statement
- XX. Answer to Quiz – Can you simplify the QUALIFY Statement
- YY. The QUALIFY Statement without Ties
- ZZ. The QUALIFY Statement with Ties
- AAA. The QUALIFY Statement with Ties Brings back Extra Rows
- BBB. Mixing Sort Order for QUALIFY Statement
- CCC. Quiz – What Caused the RANK to Reset?
- DDD. Answer to Quiz – What Caused the RANK to Reset?
- EEE. Quiz – Name those Sort Orders
- FFF. Answer to Quiz – Name those Sort Orders
- GGG. PERCENT_RANK() OVER
- HHH. PERCENT_RANK() OVER with 14 rows in Calculation

- III. PERCENT_RANK() OVER with 21 rows in Calculation
- JJJ. Quiz – What Cause the Product_ID to Reset
- KKK. Answer to Quiz – What Cause the Product_ID to Reset
- LLL. COUNT OVER for a Sequential Number
- MMM. Troubleshooting COUNT OVER
- NNN. Quiz – What caused the COUNT OVER to Reset?
- OOO. Answer to Quiz – What caused the COUNT OVER to Reset?
- PPP. The MAX OVER Command
- QQQ. MAX OVER with PARTITION BY Reset
- RRR. Troubleshooting MAX OVER
- SSS. The MIN OVER Command
- TTT. Troubleshooting MIN OVER
- UUU. The Row_Number Command
- VVV. Quiz – How did the Row_Number Reset?

XIV. Date Functions

- A. Dates are stored Internally as INTEGERS from a Formula
- B. Date, Time, and Timestamp Keywords
- C. INTEGER Date Vs. ANSIDATE is how the Date is Displayed
- D. DATEFORM
- E. Changing the DATEFORM in Client Utilities such as BTEQ
- F. Date, Time, and Timestamp Recap
- G. Timestamp Differences
- H. Troubleshooting Timestamp
- I. Add or Subtract Days from a date
- J. A Summary of Math Operations on Dates
- K. Using a Math Operation to find your Age in Years
- L. Find What Day of the week you were Born
- M. The ADD_MONTHS Command
- N. Using the ADD_MONTHS Command to Add 1-Year
- O. Using the ADD_MONTHS Command to Add 5-Years
- P. The EXTRACT Command

**Teradata
Custom for Amdocs**

Course Outline (cont'd)

- Q. EXTRACT from DATES and TIME
- R. CURRENT_DATE and Math to get Temporal Functions
- S. CAST the Date of January 1, 2011 and the Year 1800
- T. The System Calendar
- U. How to really use the Sys_Calendar.Calendar
- V. Storing Dates Internally
- W. Storing Timestamp Internally
- X. Storing Timestamp with TIME_ZONE Internally
- Y. Storing Date, Time, Timestamp with Zone Internally
- Z. Time Zones
- AA. Setting Time Zones
- BB. Seeing your Time Zone
- CC. Creating a Sample Table for Time Zone Examples
- DD. Inserting Rows in the Sample Table for Time Zone Examples
- EE. Selecting the Data from our Time Zone Table
- FF. Normalizing our Time Zone Table with a CAST
- GG. Intervals for Date, Time and Timestamp
- HH. Interval Data Types and the Bytes to Store Them
- II. The Basics of a Simple Interval
- JJ. Troubleshooting The Basics of a Simple Interval
- KK. A Date Interval Example
- LL. A Time Interval Example
- MM. A - DATE Interval Example
- NN. A Complex Time Interval Example using CAST
- OO. The OVERLAPS Command
- PP. An OVERLAPS Example that Returns No Rows
- QQ. The OVERLAPS Command using TIME
- RR. The OVERLAPS Command using a NULL Value

XV. Join Functions

- A. A two-table join using Non-ANSI Syntax
- B. Aliases and Fully Qualifying Columns
- C. A two-table join using ANSI Syntax
- D. Both Queries have the same Results and Performance
- E. Quiz – Can You Finish the Join Syntax?
- F. Answer to Quiz – Can You Finish the Join Syntax?
- G. Quiz – Can You Find the Error?
- H. Quiz – Which rows from both tables won't Return?
- I. Answer to Quiz – Which rows from both tables Won't Return?
- J. LEFT OUTER JOIN
- K. RIGHT OUTER JOIN
- L. FULL OUTER JOIN
- M. INNER JOIN with Additional AND Clause
- N. ANSI INNER JOIN with Additional AND Clause
- O. ANSI INNER JOIN with Additional WHERE Clause
- P. OUTER JOIN with Additional AND Clause
- Q. Quiz – Why is this considered an INNER JOIN?
- R. The DREADED Product Join
- S. The Horrifying Cartesian Product Join
- T. The ANSI Cartesian Join will ERROR
- U. Quiz – Do these Joins Return the Same Answer Set?
- V. Answer – Do these Joins Return the Same Answer Set?
- W. The CROSS JOIN
- X. The CROSS JOIN Answer Set
- Y. The Self Join
- Z. The Self Join with ANSI Syntax
- AA. Quiz – Will both queries bring back the same Answer Set?
- BB. How would you join these two tables?
- CC. How would you join these two tables? You can't...Yet!
- DD. An Associative Table is a Bridge that Joins Two Tables
- EE. Quiz – Can you write the 3-Table Join?

Teradata Custom for Amdocs

Course Outline (cont'd)

- FF. Answer to quiz – Can you write the 3-Table Join?
- GG. Quiz – Can you write the 3-Table Join to ANSI Syntax?
- HH. Answer – Can you write the 3-Table Join to ANSI Syntax?
- II. Quiz – Can you Place the ON Clauses at the End?
- JJ. Answer – Can you Place the ON Clauses at the End?
- KK. The 5-Table Join – Logical Insurance Model
- LL. The Nexus Query Chameleon
- MM. The 5-Table Join ANSI SQL Created by Nexus
- NN. The 5-Table Join With ON Clauses at END
- OO. The Join Tab of Nexus
- PP. The COLUMNS Tab of Nexus

Day 3 – TASM and Viewpoint

- TASM and Viewpoint
- Viewpoint
- Logging into Viewpoint
- The First Time you Login to Viewpoint
- The Add Content Menu for Monitoring
- The Add Content Menu for Monitoring
- The Add Content Menu for Tools
- The Add Content Menu for Trend Reporting
- Adding your first Content
- How the Page looks after you add your first Content
- Adding Additional Pages
- The Fundamentals of Viewpoint Pages and Portals
- Adding Multiple Portals to a Single Page
- All Portals to Their Individual Tab
- Portlet Controls
- The Rewind Control
- Query Monitor Overview
- Query Monitor Details View
- Query Monitor Actions
- My Queries
- Calendar
- Calendar Details
- Add an Event to the Calendar
- Getting a Weekly View
- System Health
- System Health Drilldown
- System Health Preferences
- ALERT
- Setting up an ALERT Configuration
- Setting up an ALERT Configuration Continued
- SQL Scratchpad
- SQL Scratchpad Query Results
- SQL Scratchpad Object Loader
- Space Usage
- Space Usage Preferences
- Investigating Space to See Dynamic Reports
- Space by Vproc (AMP)
- How to Obtain a Detailed View of Space
- Detailed View of Space
- Adding PERM Space to a Database or User
- External Content
- Lock Viewer
- Lock Viewer Lets You Configure Your View
- Canary Queries
- Multiple Systems Need Multiple Canary Query Portlets
- What Canary Queries Measure
- Remote Console
- Capacity Heatmap
- Capacity Heatmap Preferences
- Capacity Heatmap Example
- Metrics Analysis
- My Queries
- My Queries Preferences
- Today's Statistics Using System Metrics
- Today's Statistics Using Query Metrics
- Today's Statistics Using Performance Metrics
- Today's Statistics Preferences for System Metrics
- Today's Statistics Preferences for Query Metrics

Teradata Custom for Amdocs

Course Outline (cont'd)

- Today's Statistics Preferences for Performance Metrics
- Today's Statistics Preferences for Sampling Intervals
- Three Levels of Workload Management
- Pre-execution, Query Execution, and Post-execution
- What is TASM?
- Query Management compared to Workload Management
- What is the Secret Sauce for Query Management?
- The life of a Query
- What is a Workload?
- Workload Examples
- There are Four Types of Query Rules
- Common Sense Examples of Filters and Throttles
- Performance Period Examples
- The Scoop on Object Throttles
- Load Utility Throttles
- Creating Workloads
- When Creating Workloads the "WHO" is your Foundation
- After the "WHO" comes the "WHERE"
- After the "WHO" and the "WHERE" comes the "WHAT"
- Exception Actions
- When and How Teradata checks for Exceptions
- DBC.TDWMExceptionLog
- Teradata Workload Analyzer
- Teradata Workload Analyzer
- Pre-execution, Query Execution, and Post-execution
- Why use Priority Scheduler?
- The Concept of a Resource Partition
- Resource Partitions
- The Clever Idea behind Resource Partitioning
- The Brilliant Idea behind Resource Partitioning
- The Concept of Resource Partitions and Weights?
- The Concept of a Workload in a Resource Partition
- Calculating your CPU Percentage 1
- Answers to Calculating your CPU Percentage 1
- Calculating your CPU Percentage 2
- Answers to Calculating your CPU Percentage 2
- Calculating your CPU Percentage 3
- Answers to Calculating your CPU Percentage 3
- Calculating your CPU Percentage 4
- Answers to Calculating your CPU Percentage 4
- Calculating your CPU Percentage 5
- Answers to Calculating your CPU Percentage 5
- Calculating your CPU Percentage 6
- Answers to Calculating your CPU Percentage 6
- Calculating your CPU Percentage 7
- Calculating your CPU Percentage 7
- Calculating your CPU Percentage 8
- Answers to Calculating your CPU Percentage 8
- Calculating your CPU Percentage 9
- Answers to Calculating your CPU Percentage 9
- Calculating your CPU Percentage 10
- Answers to Calculating your CPU Percentage 10
- Calculating your CPU Percentage 11
- Answers to Calculating your CPU Percentage 11
- Calculating your CPU Percentage 12
- Answers to Calculating your CPU Percentage 12
- Calculating your CPU Percentage 13
- Answers to Calculating your CPU Percentage 13
- How to Configure Priority Scheduler

Teradata Custom for Amdocs

Course Outline (cont'd)

- Workload Designer
- The Three Areas of the Workload Designer
- How the Area of Workload Designer are Used
- How to Create a Ruleset
- The First Step to Creating a Ruleset
- Returning to the Workload Designer Main Screen
- A Basic Ruleset in the Working Area
- Editing a Ruleset
- The States
- States – Adding Another Planned Environment
- Editing the Name of the NewEnv Planned Environment
- Changing the Planned Environment Name to a Good Name
- Adding a New State
- A New State is Born
- Editing the Default State Name of newState
- We Now have a New State named DayState
- Drag a State (DayState) to a Planned Environment (WorkDay)
- Drag a State (DayState) to a Planned Environment (WorkDay)
- Do NOT Forget to Hit Save
- We now Need to Create a Planned Event
- Create a Planned Event to Define the Timeframes
- The Planned Event Screen
- Your Event is now in the Available Events Tab
- Drag and Drop your Event to your Planned Environment
- You have a Timeframe for the Planned Environment DayState
- Creating a Degraded Condition with our Normal Condition
- Creating a Degraded Condition with our Normal Condition
- Editing your New Condition to become a Degraded Condition
- Edit Health Conditions Screen
- Our New Edited Health Condition named Degraded
- We Now have a New Degraded Condition
- Don't Forget to Now Hit the Save Button
- Let's Define our Degraded Health when CPU hits 100%
- The Available Events Screen
- The Create Event Screen
- The Available Events Screen Now Shows our Event
- After we Drag and Drop the Event into the Degraded Events
- Drag the Base State to the Degraded/Always Environment
- Our Default of BASE is Set in the Degraded Environment
- Let's Create a New State for our Degraded Workday
- Let's Create a New State for our Degraded Workday
- Let's Create a New State for our Degraded Workday
- Drag the new state DayDegraded to our Planned Environment
- Hit Save to keep your Changes
- Limiting the Sessions
- When you First come to the Sessions Screen
- Query Session Limit – The General Tab
- Creating a Session Limit for the WorkDay when DayDegraded
- Creating a Session Limit for the WorkDay when DayDegraded
- A Successful Creation of a Session Limit on DayDegraded
- Creating Filters
- When you First come to the Filters Screen
- Creating a New Filter – The General Tab
- Creating a New Filter – The Classification Tab
- Edit Query Characteristics Criteria
- Creating a New Filter – The General Tab

**Teradata
Custom for Amdocs**

Course Outline (cont'd)

- Edit the Filter for the WorkDay when DayDegraded
- Creating a New Filter – The Classification Tab
- Utilizing Throttles
- Throttles – The Throttles Tab
- Throttles – New Throttles Tab
- Throttles – New Throttles Tab
- Edit the Throttle for the WorkDay when DayDegraded
- Throttles – New Throttles Tab
- Throttles – The Throttles Tab
- Let's Go Back and Activate our Rule Set
- To Prepare for Activation Move the Ruleset to the Ready
- The Ruleset is in the Ready and now we can Activate
- The Ruleset is in the Ready and now we can Activate
- When you First come to the Workloads Screen
- Workloads – New Workload General Tab
- Workloads – New Workload Classification Tab
- Edit the Request Source Criteria
- Workloads – New Workload Classification Tab
- Workloads – Throttles in a Workload
- Edit the DayDegraded Settings
- Workloads – Throttles are now Set for the State DayDegraded
- Workloads – Service Level Goals
- Service Level Goals for the WorkDay
- Workloads – Service Level Goals
- Workloads – Exceptions
- Workloads – The Exceptions Screen
- Workloads – The New Exceptions Screen