



Teradata Load Utilities Class

Outline

Teradata Utilities Introduction

The Teradata Utilities

Block Level Utilities

Row Level Utilities

Fast Path Inserts Using Insert/Select

Fast Path Deletes

Freespace Percent

Referential Integrity and Load Utility Solutions

Teradata has a No Primary Index Table called a NoPI Table

This is NOT Necessarily a NoPI Table

NoPI Tables Spread rows across all-AMPs Evenly

NoPI Tables used as Staging Tables for Data Loads

NoPI Table Capabilities

NoPI Table Restrictions

Why Would a NoPI Table have a Row-ID?

Collect Statistics

The Teradata Parsing Engine (Optimizer) is Cost Based

The Purpose of Collect Statistics

When Teradata Collects Statistics, it creates a Histogram

The Interval of the Collect Statistics Histogram

Histogram Quiz

Answers to Histogram Quiz

What to COLLECT STATISTICS On?

Why Collect Statistics?

How do you know if Statistics were collected on a Table?

A Huge Hint that No Statistics Have Been Collected

The Basic Syntax for COLLECT STATISTICS

COLLECT STATISTICS Examples for a better Understanding

The New Teradata V14 Way to Collect Statistics

COLLECT STATISTICS Directly From another Table

Where Does Teradata Keep the Collected Statistics?

The Official Syntax for COLLECT STATISTICS

How to Recollect STATISTICS on a Table

Teradata Always Does a Random AMP Sample

Random Sample is Kept in the Table Header in FSG Cache

Multiple Random AMP Samplings

How a Random AMP gets a Table Row count

Random AMP Estimates for NUSI Secondary Indexes

USI Random AMP Samples are Not Considered

There's No Random AMP Estimate for Non-Indexed Columns

A Summary of the PE Plan if No Statistics Were Collected

Stale Statistics Detection and Extrapolation

Extrapolation for Future Dates

How to Copy a Table with Data and the Statistics
How to Copy a Table with NO Data and the Statistics
When to COLLECT STATISTICS Using only a SAMPLE
Examples of COLLECT STATISTICS Using only a SAMPLE
Examples of COLLECT STATISTICS for V14
How to Collect Statistics on a PPI Table on the Partition
Teradata V12 and V13 Statistics Enhancements
Teradata V14 Statistics Enhancements
Teradata V14 Summary Statistics
Teradata V14 MaxValueLength
Teradata V14 MaxIntervals
Teradata V14 Sample N Percent
Teradata Statistics Wizard

Table Create and Data Types

Creating a Table with a Unique Primary Index
Creating a Table with a Non-Unique Primary Index
Creating a Table and forgetting to put in a Primary Index Clause
Creating a Set Table
Creating a Multiset Table
Creating a Set Table that won't have a Duplicate Row Check
Set Table with a Unique Constraint Eliminates the Duplicate Row Check
Creating a Table with a Unique Secondary Index
Creating a Table with a Multi-Column Primary Index
Data Types
Data Types Continued
Data Types Continued
Major Data Types and the number of Bytes they take up

Making an exact copy a Table

Making a NOT-So-Exact Copy a Table

Copying a Table with a new Default Primary Index

Troubleshooting Copying and Changing the Primary Index

Copying only specific columns of a table

Copying a Table with Data and Keeping the Statistics

Copying a Table with No Data and Statistics

Copying a table Structure with Zeroed Statistics

Creating a Table with Fallback

Creating a Table with No Fallback

Creating a Table with a Before Journal

Creating a table with a Dual Before Journal

Creating a Table with an After Journal

Creating a Table with a Dual After Journal

Creating a Table with the Journal Keyword Alone

Why use a Before Journal?

Why Use an After Journal?

Creating a Table with Customization of the Data Block Size

Creating a Table with Customization on FREESPACE Percent

Creating a QUEUE Table

Example of how a Queue Table Works

Example of how a Queue Table Works

The Concept behind Partitioning a Table

Creating a PPI Table with Simple Partitioning

Creating a PPI Table with RANGE_N Partitioning per Month

A Visual of One Year of Data with Range_N per Month

An SQL Example explaining Range_N Partitioning per Month

Creating a PPI Table with RANGE_N Partitioning per Day

Creating a PPI Table with RANGE_N Partitioning per Week

A Clever Range_N Option

Creating a PPI Table with CASE_N

NO CASE and UNKNOWN Partitions Together

Combining Older Data and Newer Data in PPI

Multi-Level Partitioning Combining Range_N and Case_N

NON-Unique Primary Indexes (NUPI) in PPI

PPI Table with a Unique Primary Index (UPI)

Tricks for Non-Unique Primary Indexes (NUPI)

Character Based PPI for RANGE_N

Character-Based PPI for CASE_N

Dates and Character-Based Multi-Level PPI

TIMESTAMP Partitioning

Using CURRENT_DATE to define a PPI

ALTER to CURRENT_DATE the next year

ALTER to CURRENT_DATE with Save

Altering a PPI Table to Add or Drop Partitions

Deleting a Partition

Deleting a Partition and saving its contents

Using the PARTITION Keyword in your SQL

SQL for RANGE_N

SQL for CASE_N

Temporary Tables

There are three types of Temporary Tables

CREATING A Derived Table

Naming the Derived Table

Aliasing the Column Names in the Derived Table

Most Derived Tables Are Used To Join To Other Tables

Multiple Ways to Alias the Columns in a Derived Table

Our Join Example with a Different Column Aliasing Style
Column Aliasing Can Default for Normal Columns
CREATING A Derived Table using the WITH Command
Our Join Example With the WITH Syntax
The Same Derived Query shown Three Different Ways
Quiz - Answer the Questions
Answer to Quiz - Answer the Questions
Clever Tricks on Aliasing Columns in a Derived Table
A Derived Table lives only for the lifetime of a single query
An Example of Two Derived Tables in a Single Query
WITH RECURSIVE Derived Table
Defining the WITH Recursive Derived Table
Looping Through the Recursive Derived Table
Looping Through a Second Time
Looping Through a Third Time
Looping Through and Adding Nothing Ends the Loop
Looping Through the WITH Recursive Derived Table
Creating a Volatile Table
You Populate a Volatile Table with an INSERT/SELECT
The Three Steps to Use a Volatile Table
Why Would You Use the ON COMMIT DELETE ROWS?
The HELP Volatile Table Command Shows your Volatiles
A Volatile Table with a Primary Index
The Joining of Two Tables Using a Volatile Table
You Can Collect Statistics on Volatile Tables
The New Teradata V14 Way to Collect Statistics
Four Examples of Creating a Volatile Table Quickly
Four Advanced Examples of Creating a Volatile Table Quickly
Creating Partitioned Primary Index (PPI) Volatile Tables

Using a Volatile Table to Get Rid of Duplicate Rows
Using a Simple Global Temporary Table
Two Brilliant Techniques for Global Temporary Tables
The Joining of Two Tables Using a Global Temporary Table
CREATING A Global Temporary Table

BTEQ – Batch Teradata Query

BTEQ – Batch Teradata Query Tool
How to Logon to BTEQ in Interactive Mode
Running Queries in BTEQ in Interactive Mode
BTEQ Commands vs BTEQ SQL Statements
WITH BY Command for Subtotals
WITH Command for a Grand Total
WITH and WITH BY Together for Subtotals and Grand Totals
How to Logon to BTEQ in a SCRIPT
Running Queries in BTEQ through a Batch Script
Running a BTEQ Batch Script through the Command Prompt
Running a BTEQ Batch Script through the Run Command
Using Nexus to Build Your BTEQ Scripts
Using Nexus to Build Your BTEQ Scripts
Using BTEQ Scripts to IMPORT Data
What Keywords Mean in a BTEQ Script
Creating a BTEQ IMPORT for a Comma Separated Value File
Four Great Examples/Ways to Run a Teradata BTEQ Script
BTEQ Export – Four types of Export Variations
Creating a BTEQ Export Script in Record Mode
Creating a BTEQ Export Script in Report Mode
The Appearance of Record Mode vs Report Mode Data
Using Report Mode to Create a Comma Separated Report

Creating a BTEQ IMPORT for a Comma Separated Value File
Using Multiple Sessions in BTEQ
BTEQ Fast Path Inserts
BTEQ Can Use Conditional Logic
Using a BTEQ Export and Setting a Limit in a UNIX System

FastLoad

FastLoad
Block Level Utility Limits
FastLoad has Two Phases
FastLoad Phase 1
FastLoad Phase 2
A Sample FastLoad Script Created by Nexus SmartScript
Executing the FastLoad Script
The Nexus SmartScript Easily Builds Your Utilities
The Nexus SmartScript FastLoad Builder
Create and Execute Your FastLoad Scripts with Nexus
FastLoad to a NoPI Table
FastLoad and CHECKPOINT
Loading Multiple Input Files with FastLoad
Valid Data Types That Can Be Used in a FastLoad
A FastLoad that Converts Data Types
A FastLoad that Uses the NULLIF Statement
FastLoad and Referential Integrity Solutions
The Output Report from FastLoad
Recovering a FastLoad that has failed
A BTEQ Export and then a FastLoad
A FastExport and then a FastLoad Needs Indicators

MultiLoad

MultiLoad

Block Level Utility Limits

MultiLoad has Five Phases

MultiLoad has IMPORT and DELETE Tasks

A Sample MultiLoad Script Created by Nexus SmartScript

Referential Integrity and Load Utility Solutions

MultiLoad That Inserts and Updates from Two Different Files

A MultiLoad Example That UPSERTs

A MultiLoad DELETE MODE Example

MultiLoad DELETE Rules

Five Formats of MultiLoad Files

A NoPI Table Does Not Work with MultiLoad

Executing a MultiLoad Script

The Output Report from MultiLoad

Host Utility Locks (HUT Locks)

Troubleshooting MultiLoad

TPump

TPump

TPump is NOT a Block Level Utility and has No Limits

Limitations of TPump

A Sample TPump Script Created by Nexus SmartScript

Executing a TPump Script

TPump Begin Load Statement Options

Five Formats of TPump Files

TPump Script with Error Treatment Options

TPump UPSERT Script

The Output Report from TPump

Did you know Tera-Tom was a world-class athlete?

FastExport

FastExport

New Rules for Block Utilities

A Sample FastExport Script Created by Nexus SmartScript

FastExport by Default places Null Indicators in Output

A Sample FastExport Script Created by Nexus SmartScript

No Spool Options with FastExport

FastExport with No Spool

FastExport that Joins Two Tables

FastExport Modes

How to Eliminate Indicators in your FastExport Script

Executing a FastExport Script

Teradata Parallel Transport (TPT)

What is TPT?

TPT Producers Create Streams and Consumers Write Them

The Four Major Operators of TPT

TPT can read from multiple source files in Parallel

TPT can have more Producers than Consumers

TPT Operators and their Equivalent Load Utility

How to Run a TPT Script

Six Syntax Rules when Creating TPT Scripts

TPT Scripts are divided into two major sections

Three Required Define Statements in the Declarative Section

The Major Keys to Building TPT Scripts

Schemas

The DDL Operator

DDL Operator Example

The SQL Selector Operator

SQL_Selector Operator Export to Delimited File (1 of 3)

SQL_Selector Operator Export to Delimited File (2 of 3)

SQL_Selector Operator Export to Delimited File (3 of 3)

Another SQL_Selector Operator Export (1 of 3)

Another SQL_Selector Operator Export (2 of 3)

Another SQL_Selector Operator Export (3 of 3)

SQL Selector Example (1 of 3)

SQL Selector Example (2 of 3)

SQL Selector Example (3 of 3)

Another SQL Selector Example (1 of 3)

Another SQL Selector Example (2 of 3)

Another SQL Selector Example (3 of 3)

The Export Operator

Export Operator to Formatted Flat File Example (1 of 3)

Export Operator to Formatted Flat File Example (2 of 3)

Export a Table to a Formatted Flat File with Indicators (3 of 3)

Deferred Schema

Export a Table to a Binary Flat File (1 of 2)

Export a Table to a Binary Flat File (2 of 2)

The Load Operator

Load from Binary File (1 of 3)

Load from Binary File (2 of 3)

Load from Binary File (3 of 3)

Data Connectors

Another Import to Table from Binary File (1 of 3)

Another Import to Table from Binary File (2 of 3)

Another Import to Table from Binary File (3 of 3)

Load Table from Flat File (1 of 2)
Load Table from Flat File (2 of 2)
Load another Table from a Flat File
Update Operator
Teradata V14.10 Extended MultiLoad Protocol (MLOADX)
Update Operator Example (1 of 3)
Update Operator Example (2 of 3)
Update Operator Example (3 of 3)
Another Update Operator Example (1 of 3)
Another Update Operator Example (2 of 3)
Another Update Operator Example (3 of 3)
Stream Operator
Stream Operator Example (1 of 3)
Stream Operator Example (2 of 3)
Stream Operator Example (3 of 3)
Another Stream Operator Example (1 of 3)
Another Stream Operator Example (2 of 3)
Another Stream Operator Example (3 of 3)
Easy Loader Example (tdload)
TPT Utility Commands
OS Command Operator
Job Variables Example
Include Statement (1 of 2)
Include Statement (2 of 2)
Operator Templates
Using Operator Templates
Operator Template Example
Moving Data from Netezza to Teradata (1 of 2)
Moving Data from Netezza to Teradata (2 of 2)

Top SQL Commands Cheat Sheet

SELECT All Columns from a Table and Sort

Select Specific Columns and Limiting the Rows

Changing your Default Database

Keywords that describe you

Select TOP Rows in a Rank Order

A Sample number of rows

Getting a Sample Percentage of rows

Find Information about a Database

Find information about a Table

Using Aggregates

Performing a Join

Performing a Join using ANSI Syntax

Using Date, Time and Timestamp

Using Date Functions

Using the System Calendar

Using the System Calendar in a Query

Formatting Data

Using Rank

Using a Derived Table

Using a Subquery

Correlated Subquery

Using Substring

Basic CASE Statement

Advanced CASE Statement

Using an Access Lock in your SQL

Collect Statistics

CREATING a Volatile Table with a Primary Index

CREATING a Volatile Table that is Partitioned (PPI)
CREATING a Volatile Table that is deleted after the Query
Finding the Typical Rows per Value for specific column
Finding out how much Space you have
How much Space you have Per AMP
Finding your Space
Finding Space Skew in Tables in a Database
Finding the Number of rows per AMP for a Column
Finding Account Information
Ordered Analytics