Greenplum SQL Class Outline

The Basics of Greenplum SQL

- Introduction
- SELECT * (All Columns) in a Table
- Fully Qualifying a Database, Schema and Table
- SELECT Specific Columns in a Table
- Commas in the Front or Back?
- Place your Commas in front for better Debugging Capabilities
- Sort the Data with the ORDER BY Keyword
- ORDER BY Defaults to Ascending
- Use the Name or the Number in your ORDER BY Statement
- Two Examples of ORDER BY using Different Techniques
- Changing the ORDER BY to Descending Order
- NULL Values sort First in Ascending Mode (Default)
- NULL Values sort Last in Descending Mode (DESC)
- Major Sort vs. Minor Sorts
- Multiple Sort Keys using Names vs. Numbers
- Sorts are Alphabetical, NOT Logical
- Using A CASE Statement to Sort Logically

WHERE Clause

The IN List is an Excellent Technique
IN List vs. OR brings the same Results
The IN List Can Use Character Data
Using a NOT IN List
Null Values in a NOT IN List Bring Back No Rows
A Technique for Handling Nulls with a NOT IN List
BETWEEN is Inclusive
NOT BETWEEN is Also Inclusive
LIKE uses Wildcards Percent '%' and Underscore '_'
LIKE command Underscore is Wildcard for one Character
ilike
LIKE Command Works Differently on Char Vs Varchar
Troubleshooting LIKE Command on Character Data
Introducing the TRIM Command
Introducing the RTRIM Command
Quiz – What Data is Left Justified and What is Right?
Numbers are Right Justified and Character Data is Left
Answer – What Data is Left Justified and What is Right?
An example of Data with Left and Right Justification
A Visual of CHARACTER Data vs. VARCHAR Data
Use the TRIM command to remove spaces on CHAR Data
Escape Character in the LIKE Command changes Wildcards
Escape Characters Turn off Wildcards in the LIKE Command
Quiz – Turn off that Wildcard
ANSWER – To Find that Wildcard
Introducing the RTRIM Command
Quiz – What Data is Left Justified and What is Right?
Answer – What Data is Left Justified and What is Right?
An example of Data with Left and Right Justification

Getting the Average Values Per Column
Average Values Per Column For all Columns in a Table
Three types of Advanced Grouping
Group By Grouping Sets
Group By Rollup
GROUP BY Rollup Result Set
GROUP BY Cube
GROUP BY CUBE Result Set
GROUP BY CUBE Result Set – check paragraph
Quiz - GROUP BY GROUPING SETS Challenge
Answer To Quiz - GROUP BY GROUPING SETS Challenge

Join Functions

Greenplum Join Quiz Answer
Redistribution
Big Table Small Table Join Strategy
Duplication of the Smaller Table across All-Distributions
If the Join Condition is the Distribution Key no Movement
Matching Rows That Are On The Same Node Naturally
Strategy 1 of 4 – The Merge Join
Quiz – Redistribute the Employees by their Dept_No
Quiz – Employees' Dept_No landed on segment with Matches
Quiz – Redistribute the Orders to the Proper segment
Answer to Redistribute the Employees by their Dept_No Quiz
Strategy 2 of 4 – The Hash Join
Strategy 3 of 4 – The Nested Join
Strategy 4 of 4 – The Product Join
A Two-Table Join Using Traditional Syntax
A two-table join using Non-ANSI Syntax with Table Alias

Date Function

- Current_Date
- Current_Date, Current_Time, and Current_Timestamp
- Current_Time vs. LocalTime With Precision
- Local_Time and Local_Timestamp With Precision
- Now() and Timeofday() Functions
- Adding A Week to a Date
- Add or Subtract Days from a date
- Formatting Dates and Dollar Amounts
- The EXTRACT Command
- EXTRACT from DATES and TIME
- EXTRACT Command on the Century
- EXTRACT Command for the Decade, DOW and DOY
- EXTRACT Microseconds, Milliseconds and Millennium
- EXTRACT of the Month on Aggregate Queries
- Date_part Command
- Date_Trunc Command With Time
- Date_Trunc Command With Dates
- The AGE Command
- AGE Challenge
- AGE Challenge Results
- Epoch
- Using Intervals
- More Interval Examples
- Interval Arithmetic Results
- A Complex Time Interval example using CAST
- The OVERLAPS Command
- An OVERLAPS example that Returns No Rows
- The OVERLAPS Command using TIME

Using Both CAST and CONVERT in Literal Values
A Better Technique for YEAR, MONTH, and DAY Functions

Conversions and Formatting

Postgres Conversion Functions
Postgres Conversion Function Templates
Postgres Conversion Function Templates Continued
To_Char command Examples
Formatting A Date with To_Char
Formatting A Date With To_Char Continued
To_Number
The to_number syntax, and postgres parameters and their explanation are listed above. Not all above work on Greenplum
To_Number Examples
To_Date
To_Timestamp

Sub-query Functions

An IN List is much like a Subquery
An IN List Never has Duplicates – Just like a Subquery
An IN List Ignores Duplicates
The Subquery
The Three Steps of How a Basic Subquery Works
These are Equivalent Queries
The Final Answer Set from the Subquery
Quiz- Answer the Difficult Question
Answer to Quiz- Answer the Difficult Question
Should you use a Subquery of a Join?
Quiz- Write the Subquery

Answer to Quiz- Write the Subquery
Quiz- Write the More Difficult Subquery
Answer to Quiz- Write the More Difficult Subquery
Quiz – Write the Extreme Subquery
Answer to Quiz – Write the Extreme Subquery
Quiz- Write the Subquery with an Aggregate
Answer to Quiz- Write the Subquery with an Aggregate
Quiz- Write the Correlated Subquery
Answer to Quiz- Write the Correlated Subquery
The Basics of a Correlated Subquery
The Top Query always runs first in a Correlated Subquery
Correlated Subquery Example vs. a Join with a Derived Table
Quiz- A Second Chance to Write a Correlated Subquery
Answer - A Second Chance to Write a Correlated Subquery
Quiz- A Third Chance to Write a Correlated Subquery
Answer - A Third Chance to Write a Correlated Subquery
Quiz- Last Chance To Write a Correlated Subquery
Answer – Last Chance to Write a Correlated Subquery
Quiz – Write the Extreme Correlated Subquery
Answer To Quiz – Write the Extreme Correlated Subquery
Quiz- Write the NOT Subquery
Answer to Quiz- Write the NOT Subquery
Quiz- Write the Subquery using a WHERE Clause
Answer - Write the Subquery using a WHERE Clause
Quiz- Write the Subquery with Two Parameters
Answer to Quiz- Write the Subquery with Two Parameters
More on how the Double Parameter Subquery Works
Quiz – Write the Triple Subquery
Answer to Quiz – Write the Triple Subquery

Quiz – How many rows return on a NOT IN with a NULL?
Answer – How many rows return on a NOT IN with a NULL?
How to handle a NOT IN with Potential NULL Values
IN is equivalent to =ANY
Using a Correlated Exists
How a Correlated Exists matches up
The Correlated NOT Exists
The Correlated NOT Exists Answer Set
Quiz – How many rows come back from this NOT Exists?
Answer – How many rows come back from this NOT Exists?

OLAP Functions

CSUM
CSUM – The Sort Explained
CSUM – Rows Unbounded Preceding Explained
CSUM – Making Sense of the Data
CSUM – Making Even More Sense of the Data
CSUM – The Major and Minor Sort Key(s)
The ANSI CSUM – Getting a Sequential Number
Troubleshooting The ANSI OLAP on a GROUP BY
Reset with a PARTITION BY Statement
PARTITION BY only Resets a Single OLAP not ALL of them
Moving SUM
ANSI Moving Window is Current Row and Preceding n Rows
How ANSI Moving SUM Handles the Sort
Quiz – How is that Total Calculated?
Answer to Quiz – How is that Total Calculated?
Moving SUM every 3-rows Vs a Continuous Average
Partition By Resets an ANSI OLAP

Both the Greenplum Moving Average and ANSI Version
Moving Average
The Moving Window is Current Row and Preceding
How Moving Average Handles the Sort
Quiz – How is that Total Calculated?
Answer to Quiz – How is that Total Calculated?
Quiz – How is that 4th Row Calculated?
Answer to Quiz – How is that 4th Row Calculated?
Moving Average every 3-rows Vs a Continuous Average
Partition By Resets an ANSI OLAP
Moving Difference using ANSI Syntax with Partition By
RANK Defaults to Ascending Order
Getting RANK to Sort in DESC Order
RANK() OVER and PARTITION BY
RANK and DENSE RANK
PERCENT_RANK() OVER
PERCENT_RANK() OVER with 14 rows in Calculation
PERCENT_RANK() OVER with 21 rows in Calculation
Quiz – What Causes the Product_ID to Reset?
Answer to Quiz – What Cause the Product_ID to Reset?
COUNT OVER for a Sequential Number
Troubleshooting COUNT OVER
Quiz – What caused the COUNT OVER to Reset?
Answer to Quiz – What caused the COUNT OVER to Reset?
The MAX OVER Command
MAX OVER with PARTITION BY Reset
Troubleshooting MAX OVER
The MIN OVER Command
Troubleshooting MIN OVER

Finding a Value of a Column in the Next Row with MIN
Quiz – Fill in the Blank
Answer – Fill in the Blank
The Row_Number Command
Using a Derived Table and Row_Number
Quiz – How did the Row_Number Reset?
Answer – How did the Row_Number Reset?
Ordered Analytics OVER
CURRENT ROW AND UNBOUNDED FOLLOWING
Different Windowing Options
The CSUM For Each Product_Id and the Next Start Date
How Ntile Works
Ntile
Ntile Continued
Ntile Percentile
Another Ntile example
Using Tertiles (Partitions of Four)
NTILE
NTILE Using a Value of 10
NTILE With a Partition
Using FIRST_VALUE
FIRST_VALUE
FIRST_VALUE After Sorting by the Highest Value
FIRST_VALUE with Partitioning
Using LAST_VALUE
LAST_VALUE
Using LEAD
Using LEAD With and Offset of 2
LEAD

Temporary Tables

Substrings and Positioning Functions

Answer - Valued Case Statement
Quiz - Searched Case Statement
Answer - Searched Case Statement
The CASE Challenge
The CASE Challenge Answer
Combining Searched Case and Valued Case
A Trick for getting a Horizontal Case
Nested Case

Set Operators Functions

Rules of Set Operators
Rules of Set Operators
INTERSECT Explained Logically
INTERSECT Explained Logically
UNION Explained Logically
UNION Explained Logically
UNION ALL Explained Logically
UNION ALL Explained Logically
EXCEPT Explained Logically
EXCEPT Explained Logically
An Equal Amount of Columns in both SELECT List
Columns in the SELECT list should be from the same Domain
The Top Query handles all Aliases
The Bottom Query does the ORDER BY (a Number)
Great Trick: Place your Set Operator in a Derived Table
UNION Vs UNION ALL
Using UNION ALL and Literals
A Great example of how EXCEPT works
Quiz – Build that Query

Statistical Aggregate Functions

A VARIANCE Example
The CORR Function
A CORR Example
Another CORR Example so you can Compare
The COVAR_POP Function
A COVAR_POP Example
Another COVAR_POP Example so you can Compare
The COVAR_SAMP Function
A COVAR_SAMP Example
Another COVAR_SAMP Example so you can Compare
The REGR_INTERCEPT Function
A REGR_INTERCEPT Example
Another REGR_INTERCEPT Example so you can Compare
The REGR_SLOPE Function
A REGR_SLOPE Example
Another REGR_SLOPE Example so you can Compare
The REGR_AVGX Function
A REGR_AVGX Example
Another REGR_AVGX Example so you can Compare
The REGR_AVGY Function
A REGR_AVGY Example
Another COVAR_POP Example so you can Compare
The REGR_COUNT Function
A REGR_COUNT Example
The REGR_R2 Function
A REGR_R2 Example
The REGR_SXX Function
A REGR_SXX Example
The REGR_SXY Function