ORACLE COST-BASED OPTIMIZER ADVANCED

Randolf Geist
http://oracle-randolf.blogspot.com
info@sqltools-plusplus.org

ABOUT ME

- Independent consultant
  - Available for consulting
  - In-house workshops
    - Cost-Based Optimizer
    - Performance By Design
  - Performance Troubleshooting
- Oracle ACE Director
- Member of OakTable Network

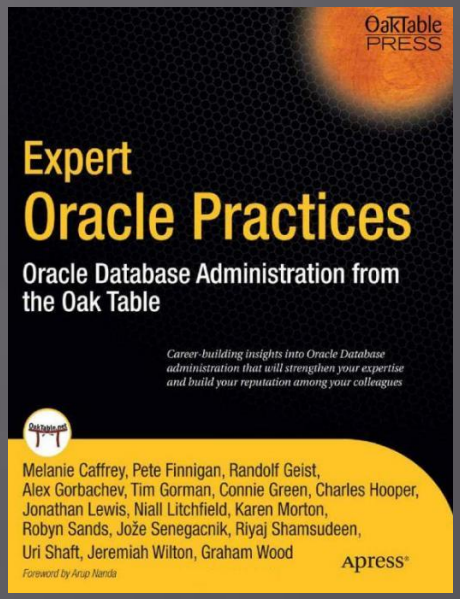

OPTIMIZER BASICS

- Three main questions you should ask when looking for an efficient execution plan:
  - How much data? How many rows / volume?
  - How scattered / clustered is the data?
  - Caching?

=> Know your data!

OPTIMIZER BASICS

- Why are these questions so important?
  - Two main strategies:
    - One “Big Job”
      => How much data, volume?
    - Few/many “Small Jobs”
      => How many times / rows?
      => Effort per iteration? Clustering / Caching

OPTIMIZER BASICS

- Optimizer's cost estimate is based on:
  - How much data? How many rows / volume?
  - How scattered / clustered? (partially)
  - (Caching?) Not at all

BASICS' SUMMARY

- Cardinality and Clustering determine whether the "Big Job" or "Small Job" strategy should be preferred
- If the optimizer gets these estimates right, the resulting execution plan will be efficient within the boundaries of the given access paths
- Know your data and business questions

AGENDA

- Clustering Factor
- Statistics / Histograms
- Datatype issues

HOW SCATTERED / CLUSTERED?

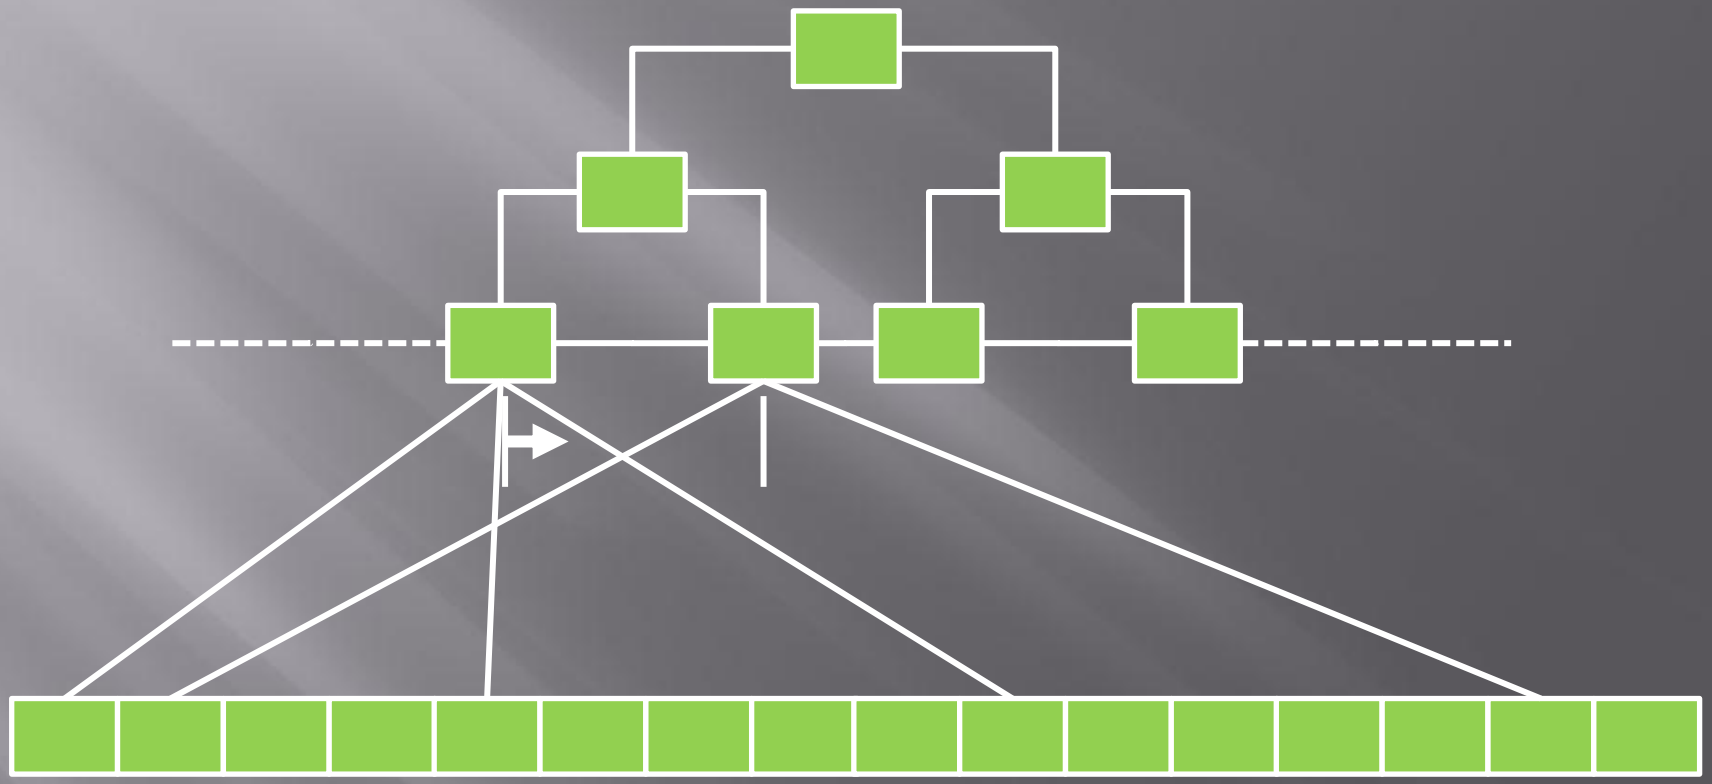

1,000 rows => visit 1,000 table blocks: 1,000 * 5ms = 5 s

HOW SCATTERED / CLUSTERED?

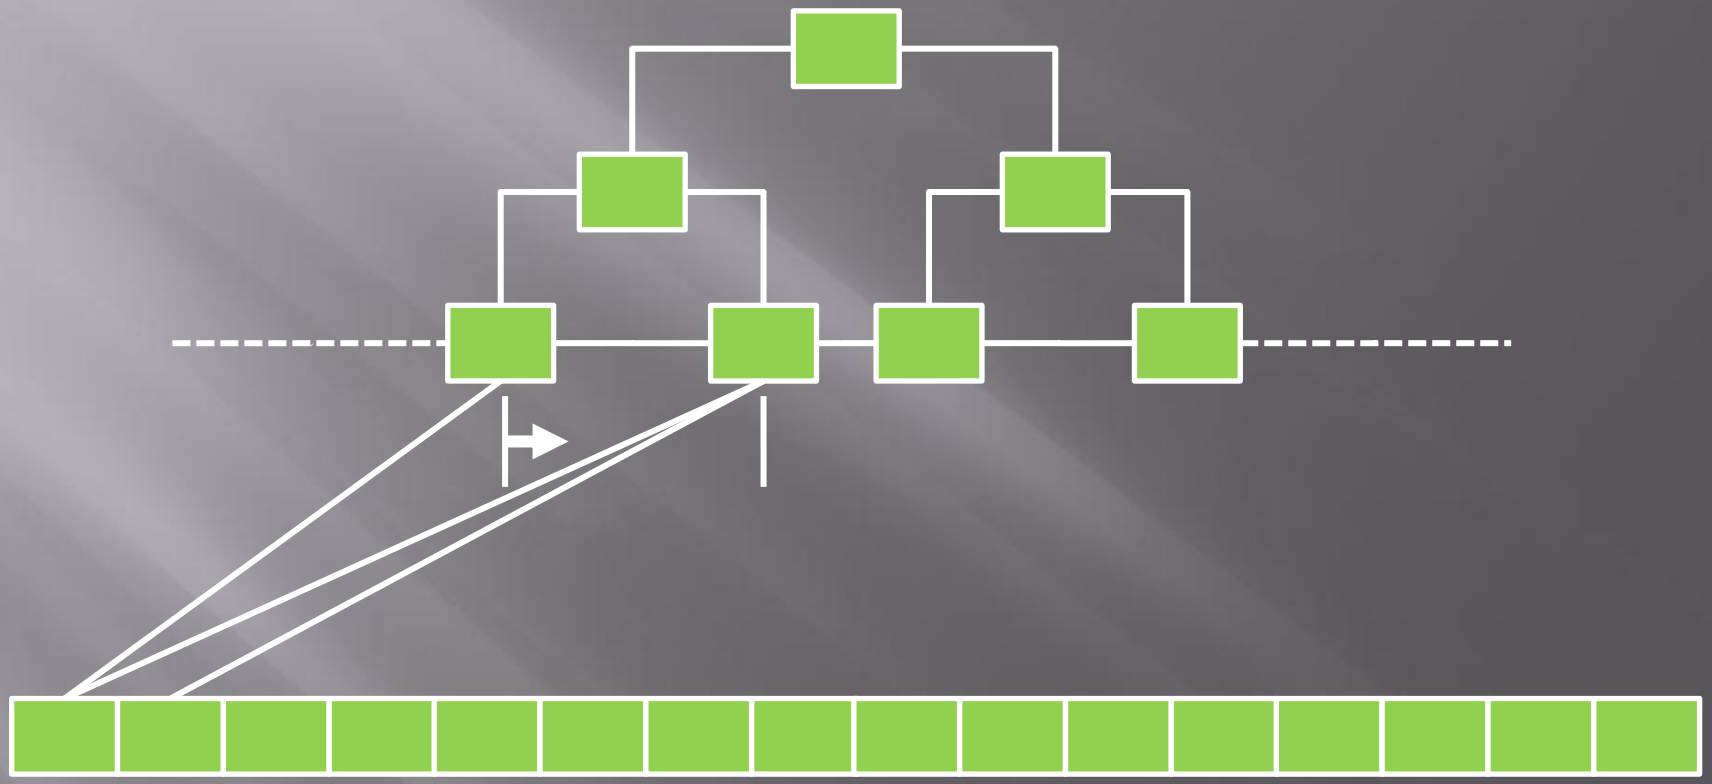

1,000 rows => visit 10 table blocks: 10 * 5ms = 50 ms

HOW SCATTERED / CLUSTERED?

- There is only a single measure of clustering in Oracle:
  The index clustering factor
- The index clustering factor is represented by a single value
- The logic measuring the clustering factor by default does not cater for data clustered across few blocks (ASSM!)

HOW SCATTERED / CLUSTERED?

- Challenges
  - Getting the index clustering factor right
  - There are various reasons why the index clustering factor measured by Oracle might not be representative
    - Multiple freelists / freelist groups (MSSM)
    - ASSM
    - Partitioning
    - SHRINK SPACE effects

HOW SCATTERED / CLUSTERED?

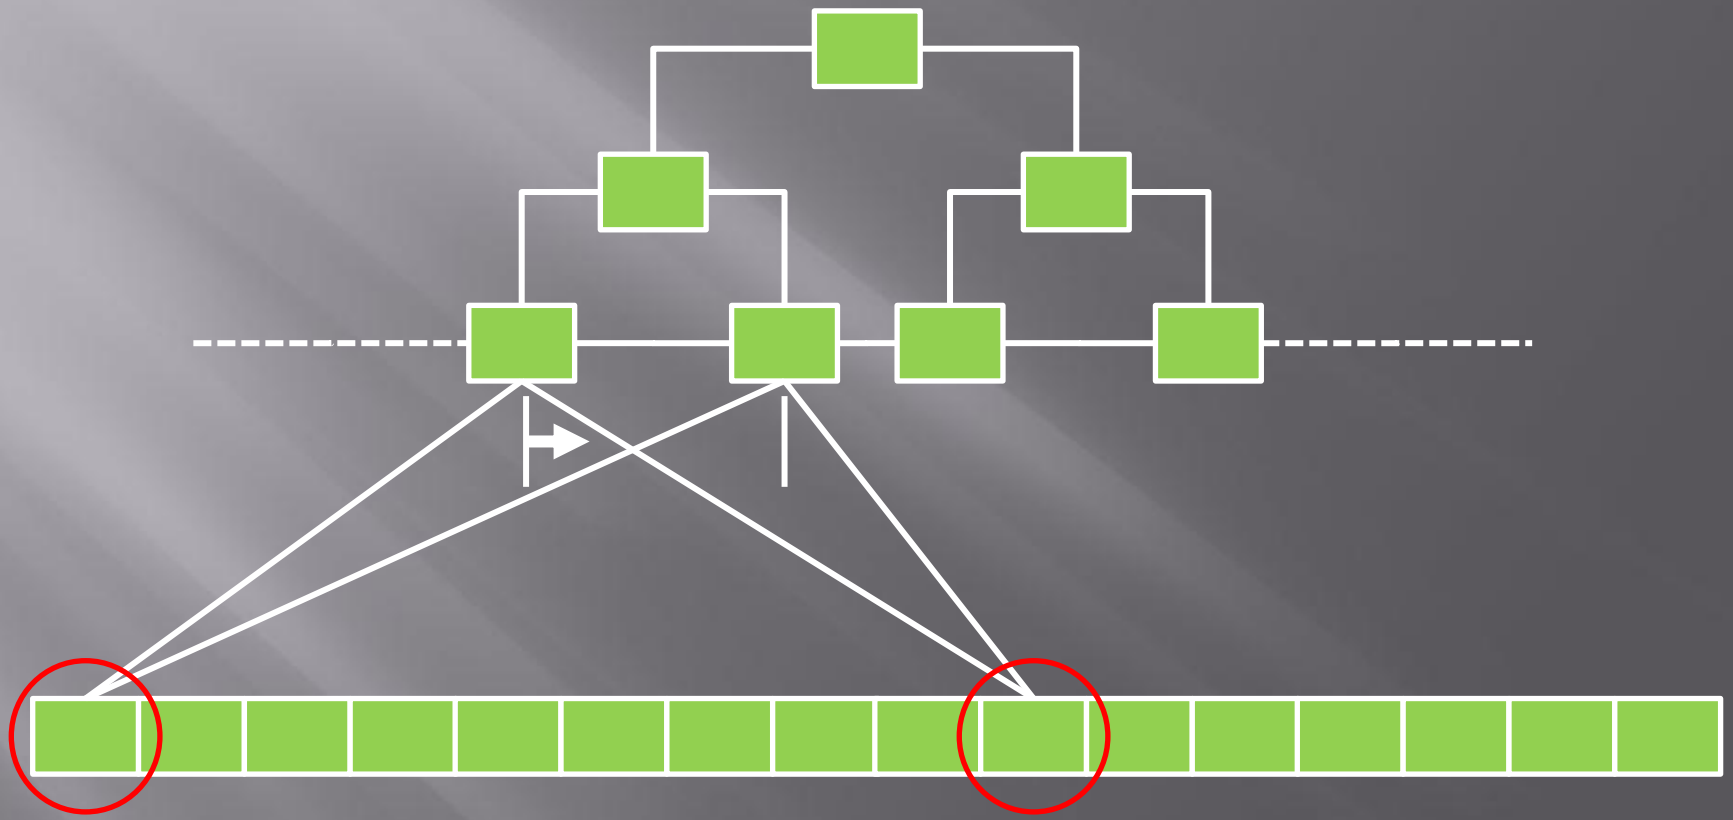

Re-visiting the same recent table blocks

STATISTICS

- Don't use ANALYZE … COMPUTE / ESTIMATE STATISTICS anymore
- Basic Statistics:

  - Table statistics: Blocks, Rows, Avg Row Len
  Nothing to configure there, always generated

  - Basic Column Statistics: Low / High Value, Num Distinct, Num Nulls
  => Controlled via METHOD_OPT option of DBMS_STATS.GATHER_TABLE_STATS

STATISTICS

Controlling column statistics via METHOD_OPT

- If you see FOR ALL INDEXED COLUMNS [SIZE > 1]: Question it! Only applicable if the author really knows what he/she is doing! => Without basic column statistics Optimizer is resorting to hard coded defaults!
- Default in previous releases:
  FOR ALL COLUMNS SIZE 1: Basic column statistics for all columns, no histograms
- Default from 10g on:
  FOR ALL COLUMNS SIZE AUTO: Basic column statistics for all columns, histograms if Oracle determines so

HISTOGRAMS

- Basic column statistics get generated along with table statistics in a single pass (almost)
- Each histogram requires a separate pass
- Therefore Oracle resorts to aggressive sampling if allowed => AUTO_SAMPLE_SIZE
- This limits the quality of histograms and their significance

HISTOGRAMS

- Limited resolution of 255 value pairs maximum
- Less than 255 distinct column values => Frequency Histogram
- More than 255 distinct column values => Height Balanced Histogram
- Height Balanced is always a sampling of data, even when computing statistics!

FREQUENCY HISTOGRAMS

- SIZE AUTO generates Frequency Histograms if a column gets used as a predicate and it has less than 255 distinct values
- Major change in behaviour of histograms introduced in 10.2.0.4 / 11g
- Be aware of new "value not found in Frequency Histogram" behaviour
- Be aware of edge case of very popular / unpopular values

HEIGHT BALANCED HISTOGRAMS

SELECT SKEWED_NUMBER FROM T ORDER BY SKEWED_NUMBER

10,000,000 rows

100 popular values
with 70,000 occurrences

250 buckets each covering
40,000 rows (compute)

250 buckets each covering
approx. 22/23 rows (estimate)

HEIGHT BALANCED HISTOGRAMS

HEIGHT BALANCED HISTOGRAMS

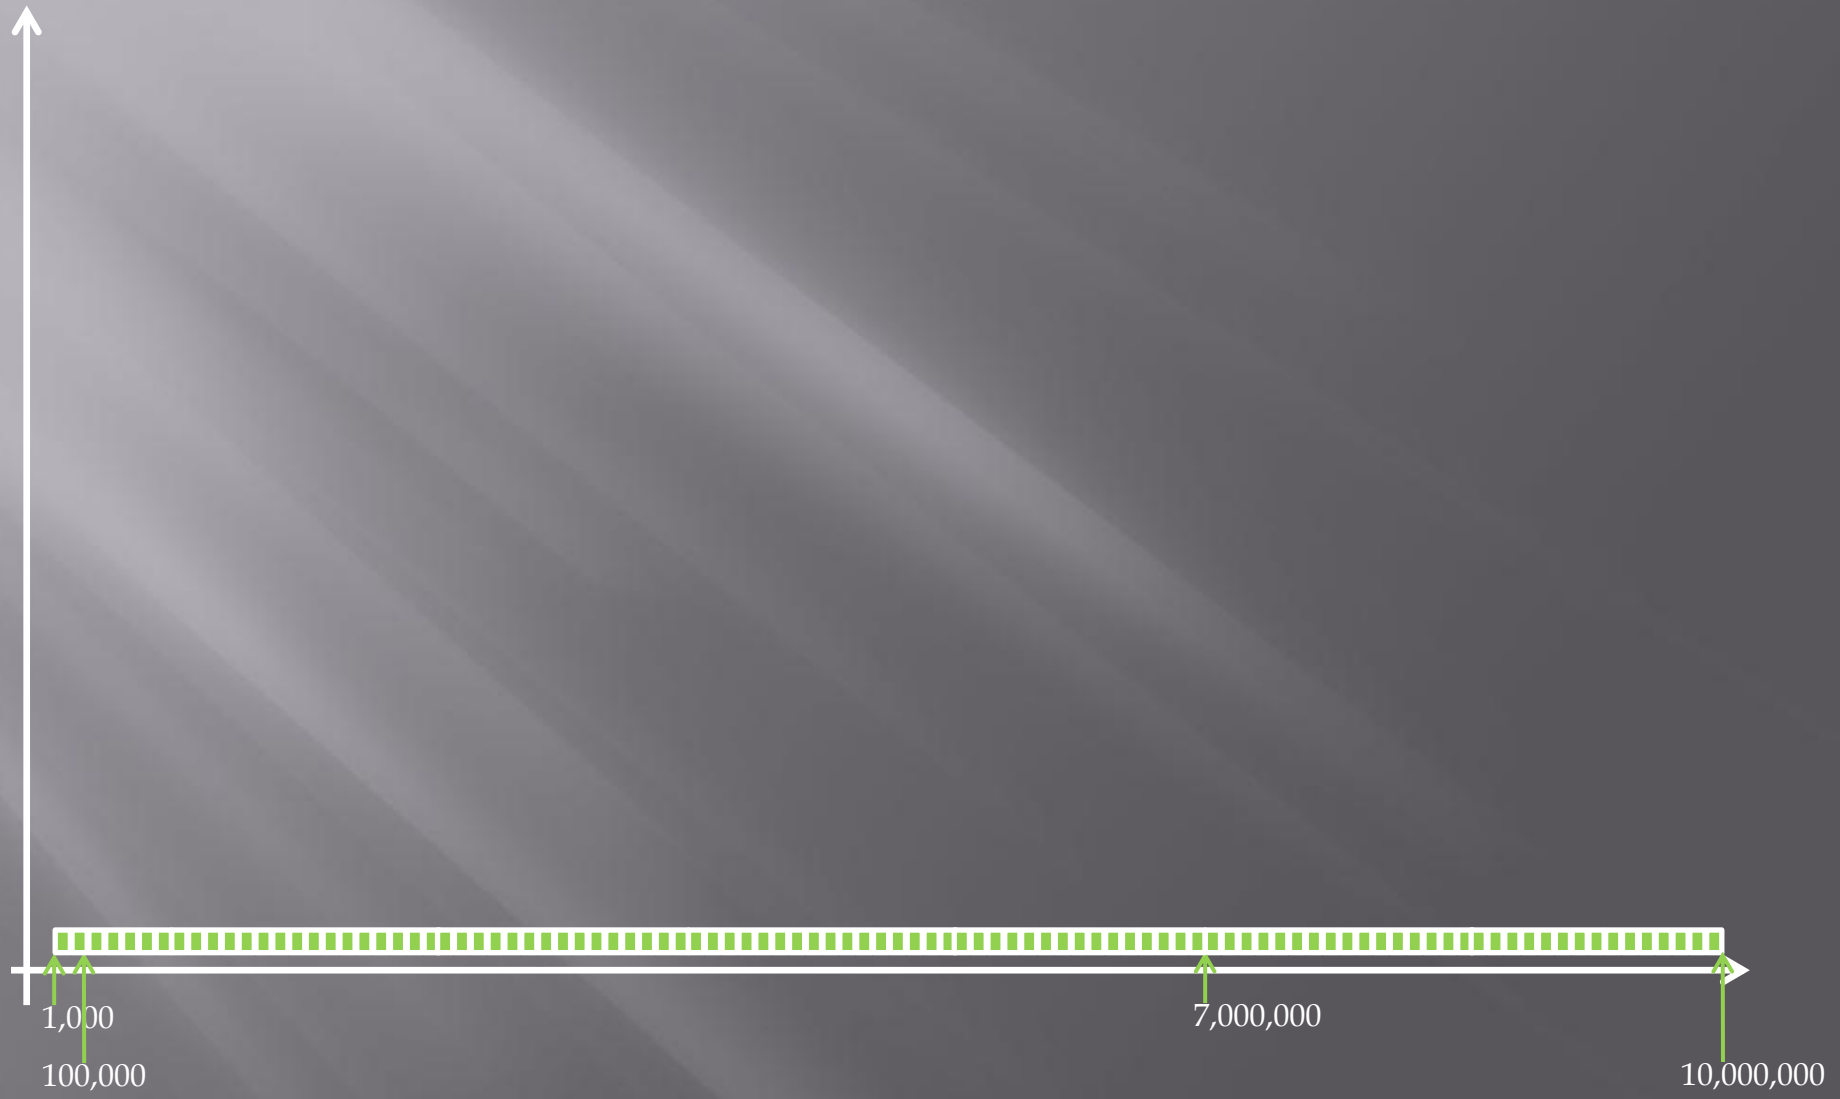

SUMMARY

- Check the correctness of the clustering factor for your critical indexes
- Oracle does not know the questions you ask about the data
- You may want to use FOR ALL COLUMNS SIZE 1 as default and only generate histograms where really necessary
- You may get better results with the old histogram behaviour, but not always

SUMMARY

- There are data patterns that don't work well with histograms when generated via Oracle
- => You may need to manually generate histograms using DBMS_STATS.SET_COLUMN_STATS for critical columns
- Don't forget about Dynamic Sampling / Function Based Indexes / Virtual Columns / Extended Statistics
- Know your data and business questions!

QUESTIONS & ANSWERS

Q & A