

## **Introduction**

- Goals
- Outline
- Q&A (quick questions before we start)

## **The Modern Data Warehouse**

- Key Analytic Data Challenges of the 2020's
- Evolution of the Data Warehouse
- Scale, Complexity & Risk
- Operating Cost in the Cloud
- Data Warehousing in the Cloud

## **Cloud Architecture for the Data Warehouse**

- Provider Landscape
- Regions and Zones
- Pricing
- Storage Types and Tradeoffs
- Multi-Cloud
- Trends

## **Architecture of the Modern Data Warehouse Platform , Part 1**

- Parallelism
- Elasticity
- Partitioning/Clustering/Indexing

## **Architecture of the Modern Data Warehouse Platform, Part 2**

- Query Processing
- Joins
- Views, Materialized Views
- Concurrency, Workload Management

## **Key Cloud Data Warehouse Platforms**

- AWS Redshift
- Cloudera Data Warehouse
- Data Bricks
- Google BigQuery

- Microsoft Azure Synapse Analytics
- Oracle Autonomous Data Warehouse (ADW)
- Snowflake
- Teradata
- Others of Interest

**Summary Product Matrix**

**Tips on Choosing Your Platform**

**Q&A**