Analyze in the Right Place

…for the right price

Time-limited pricing for analyzing data in your data lake.

Many Vertica customers have growing data lakes and ever increasing data volumes in Apache Hadoop (HDFS) or AWS S3. Hadoop and S3 help companies store massive volumes of data at a low cost and make for operational efficiencies.

However, performing analytics on the data lake can often require that you move data into your data analytics engine for querying. With its external tables features, Vertica can analyze that data in-place without the need to move it or subsample it.

Using our unified analytics engine, users can perform analytics on data no matter where it lives—Hadoop, native Vertica, or in the clouds. There is no need for data movement and no need to copy the data. Data is analyzed in-place by defining it as an external table and using that table just as you would use a Vertica ROS file. Vertica can quickly and efficiently query ORC and Parquet files for fast Hadoop data analytics without moving the data. The analytics provided is not only SQL 99 compliant, but allows for predictive analytics using Vertica's in-database machine learning capabilities.

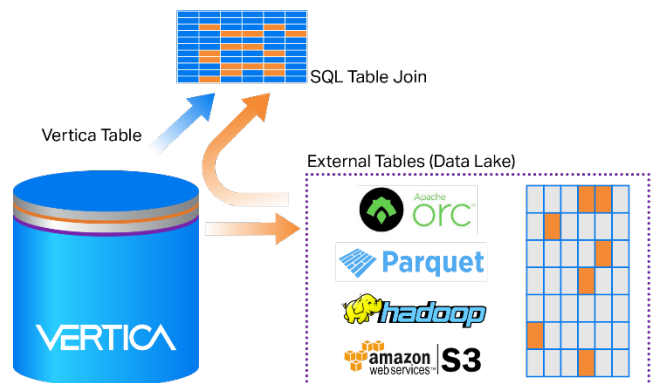### Why External Tables are so crucial to your analytics

There are four key reasons why External Tables matter to the entire analytics community within your company:

Data Science – Now you can use the Vertica engine to explore data. Vertica will run 100% of your SQL queries and supports your visualization tools, SQL, R, and Python. Even when data volumes are too massive and the demand is too overwhelming for many concurrent queries on standard Hadoop, Vertica can meet the challenge.

JOINing External Data – You can gain additional insight by creating JOINs between your Vertica data warehouse and data that is sitting untapped in Hadoop. For example, users can leveraging Web logs to gain additional customer insight, use sensor and IOT data that is sitting externally for pre-emptive service, or track the success of marketing programs by joining data with your data warehouse.

Training the Machine Learning Models – A common practice is to move data, or subsets of data, from the data warehouse into Apache Spark to train your machine learning models. However, using Vertica's External Tables feature and built-in machine learning capabilities, you can train your models without having to move the data. You can use more amounts of data and get more precision on your models.

Information Lifecycle Management – After your hot data cools off, a common practice is to move your data off to lower-cost storage in support of information lifecycle management. Now, with our Parquet WRITE capabilities, you can move cooler data out of Vertica into Hadoop or the Clouds and still analyze it with Vertica using External Tables.

The Offer

We are excited to offer a highly discounted price per/TB for using External Tables in your data lake in Vertica. This discount applies to data stored in Parquet or ORC format when stored in Hadoop or S3. This offer is specifically for new or existing Vertica Enterprise license holders. The per/TB list price on a perpetual license for externally stored data is now 25% of data stored in native Vertica format. Data stored in External Tables is naturally compressed and the pricing is measured against the compressed footprint. **Learn more at www.vertica.com or contact your Vertica sales representative to take advantage of this special, time-limited offer.**