



Hewlett Packard
Enterprise



Vertica and Spark

Connecting Computation and Data

Edward Ma

Motivation

- **Apache Spark** is a popular computation engine
 - Excellent integration with Hadoop ecosystem
 - Has a generalized multistage in-memory computational paradigm
 - Scheduler assumes many independent, non-coordinated tasks
- Data-Management Issues
 - Lack of adequate persistent data management capability
 - Spark SQL lacks enterprise-grade data management (isolation, integrity)
 - Spark has APIs for external data integration



Vertica

– Fast, MPP Analytics Database

– Enterprise-grade data management (isolation, integrity, etc.)

– An ideal data source for Spark ...



Spark External Data Source API (Spark SQL)

```
val df = sqlContext.read.format("com.vertica.DefaultSource").load()
```

```
df.select("name", "age").filter("age" < 50)
```

Permits Filters and Projections to be pushed down to the source




Our Approach

- **Addresses challenges of Spark**
 - Spark scheduler assumes many independent, non-coordinated tasks
 - Issue direct, mutually-exclusive queries for different data ranges
 - Queries can be reissued and recomputed independently of others
- Filter and Projection pushdowns can be trivially parallelized to each individual query.
- **Key Innovation: Leverage Data Locality**


Vertica Table Data Segmentation: Hash Ring

- node0
- node1
- node2
- node3



Vertica Table Data Segmentation: Hash Ring

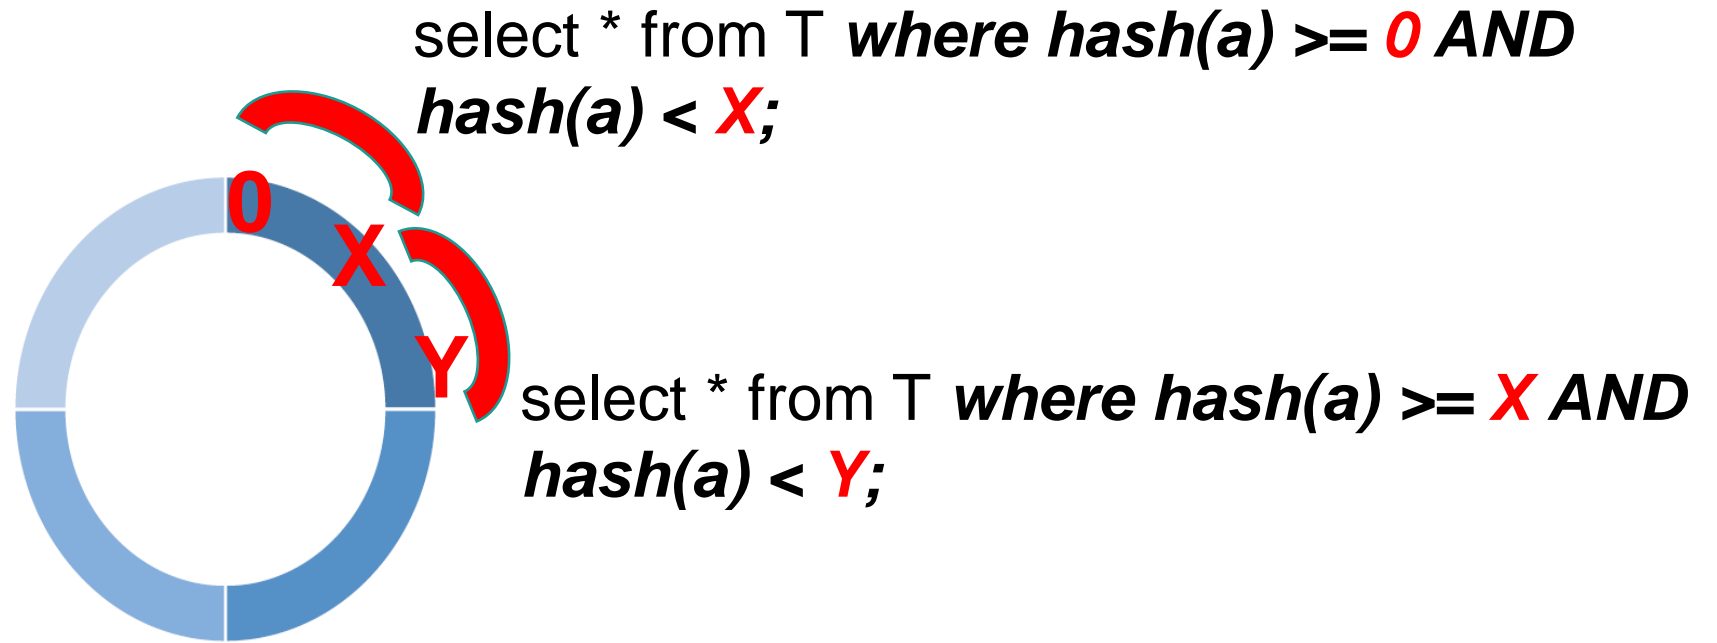
- node0
- node1
- node2
- node3



Spark Partition

Vertica Table Data Segmentation: Hash Ring

- node0
- node1
- node2
- node3



 Spark Partition

Optimization: Make query plans local with a hint

- A query hint that indicates the query may be only executed on initiator node

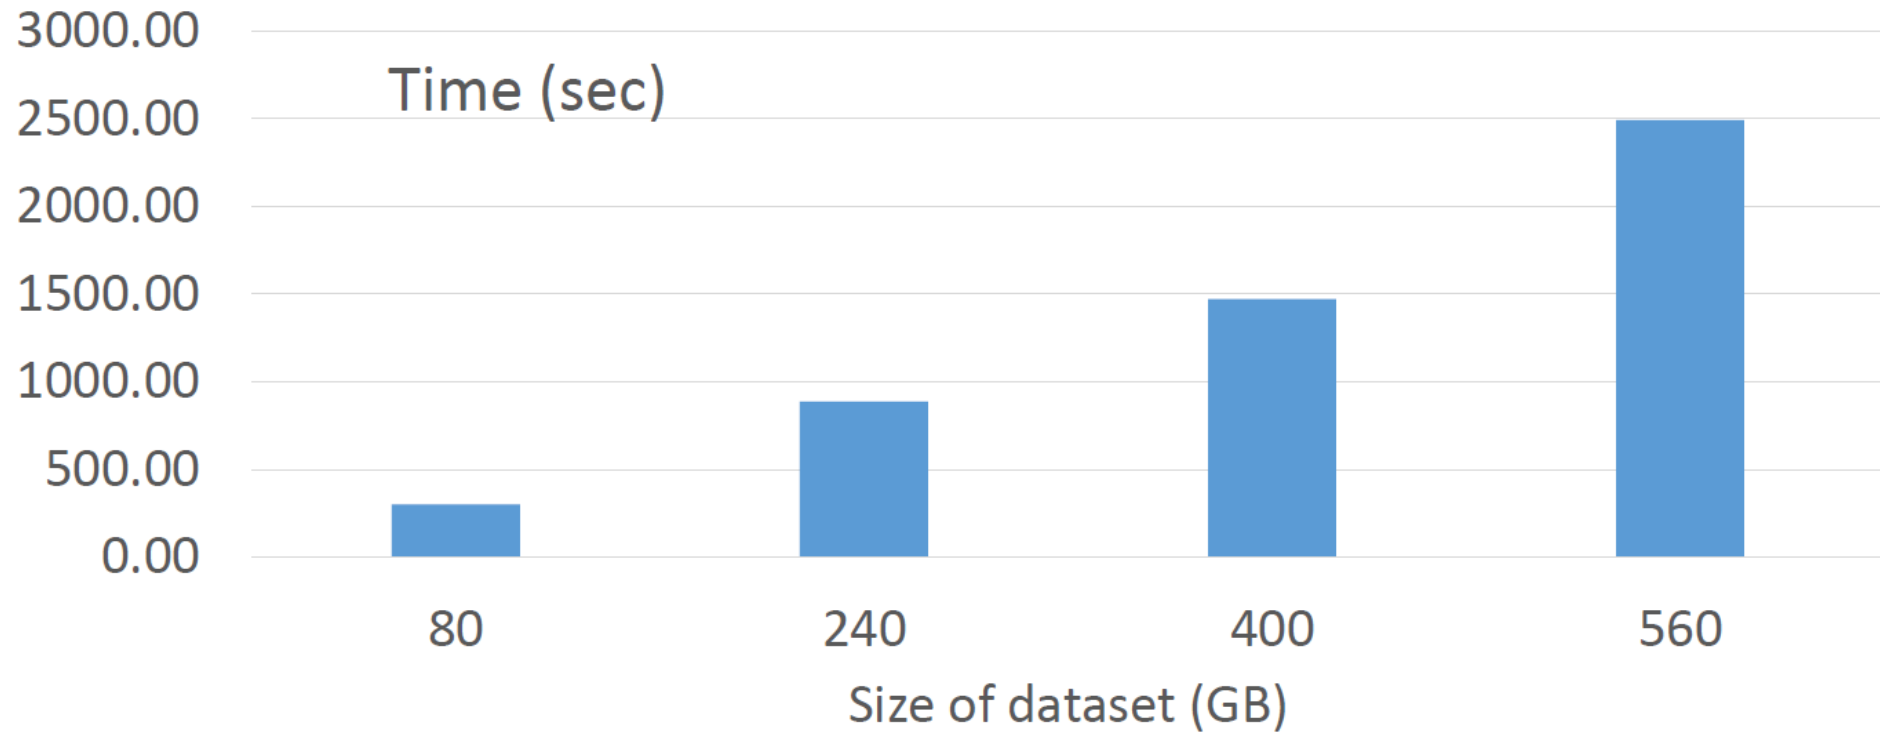
```
select /*+initiator_local*/ * from T where hash(a) >= 0 AND hash(a) < X;
```

Benefits of hint and data locality:

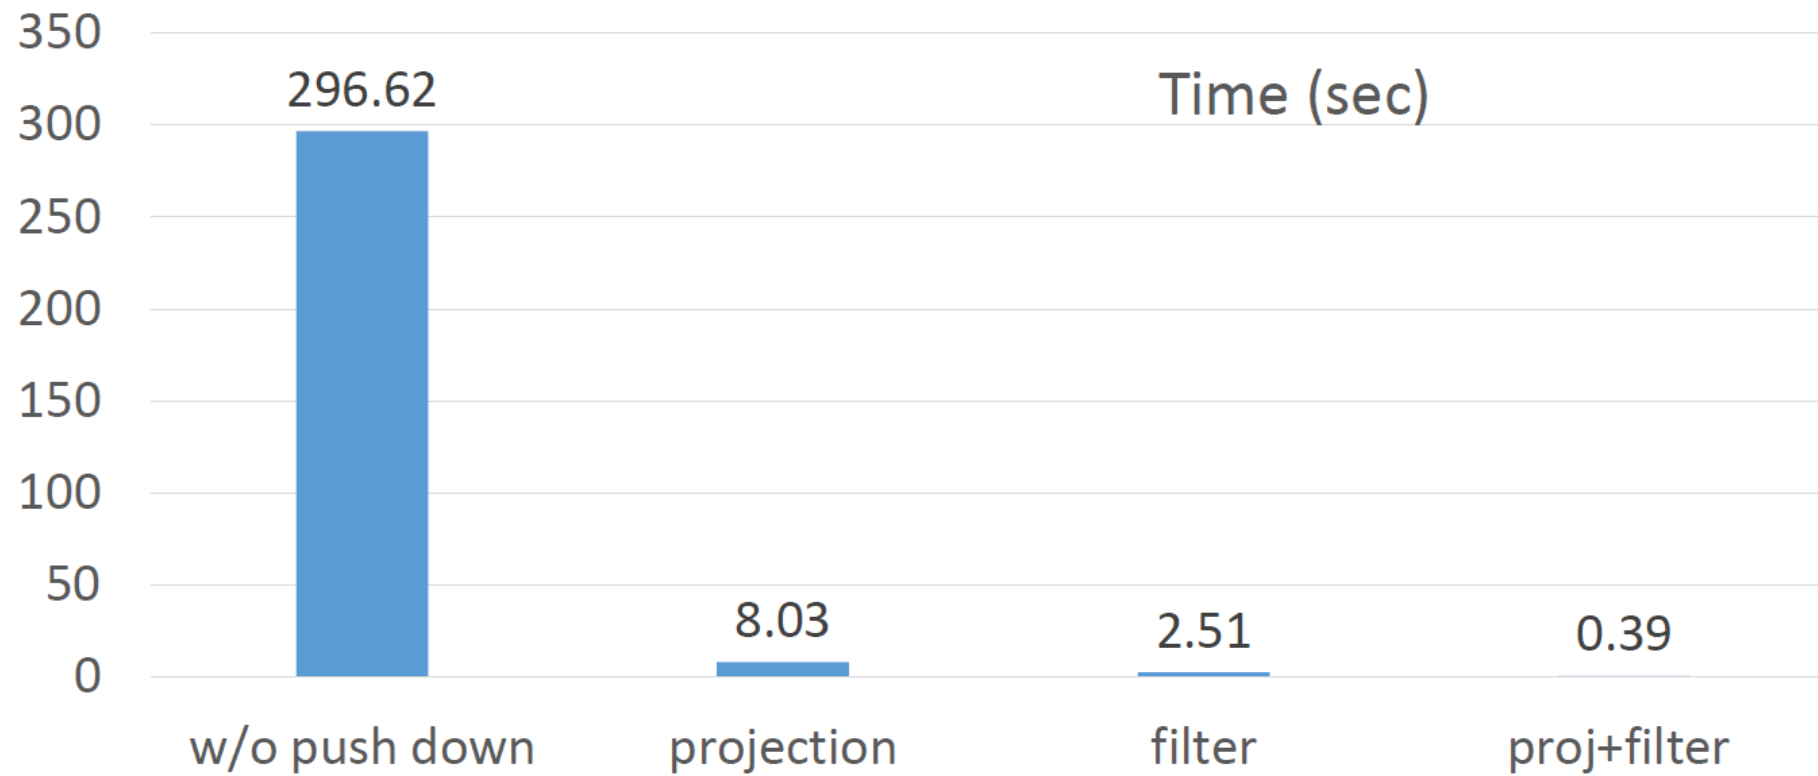
- No traffic between Vertica nodes → no wasted internode network bandwidth
- Other executor nodes do not need to run anything → no wasted CPU resource



Throughput and Data Scalability



Performance Impact of Pushdowns





Hewlett Packard
Enterprise

Thank you

ema@hpe.com