Presto at Wayfair

Vinay Narayana

www.linkedin.com/in/vinaynarayana/

@nvinay26

1. Problem Statement

2. Why Presto ?

3. Presto at Wayfair

   Deployment
   Adoption
   Performance
   Monitoring

4. What's Next

Problem Statement

Remedies

1. Optimize Hive queries
2. Set up queues to prioritize batch jobs
3. Throttle users to 2 ad-hoc hive queries
4. Move jobs from Hive to Spark
5. Conduct SME training session for both Hive and Spark

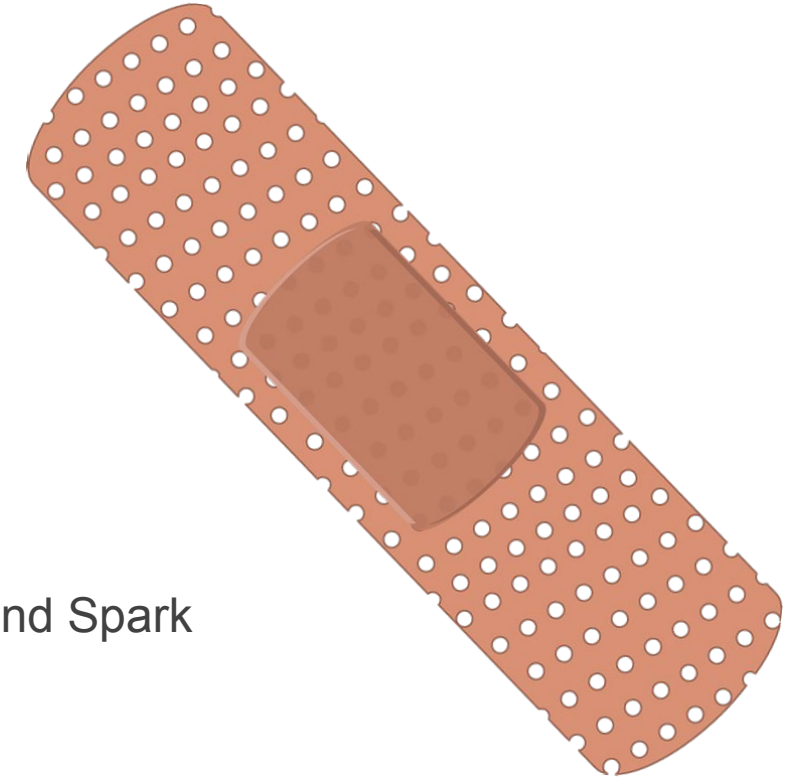

Why Presto ?

- It’s VERY fast!
- ANSI SQL Support
- Presto can run separately from the storage HDFS cluster making it great for interactive queries
- Single SQL query to access, combine and analyze data from multiple data sources (unlike Impala)
- Presto is easier to understand and use versus Spark

Presto At Wayfair

Presto Ad Hoc Cluster

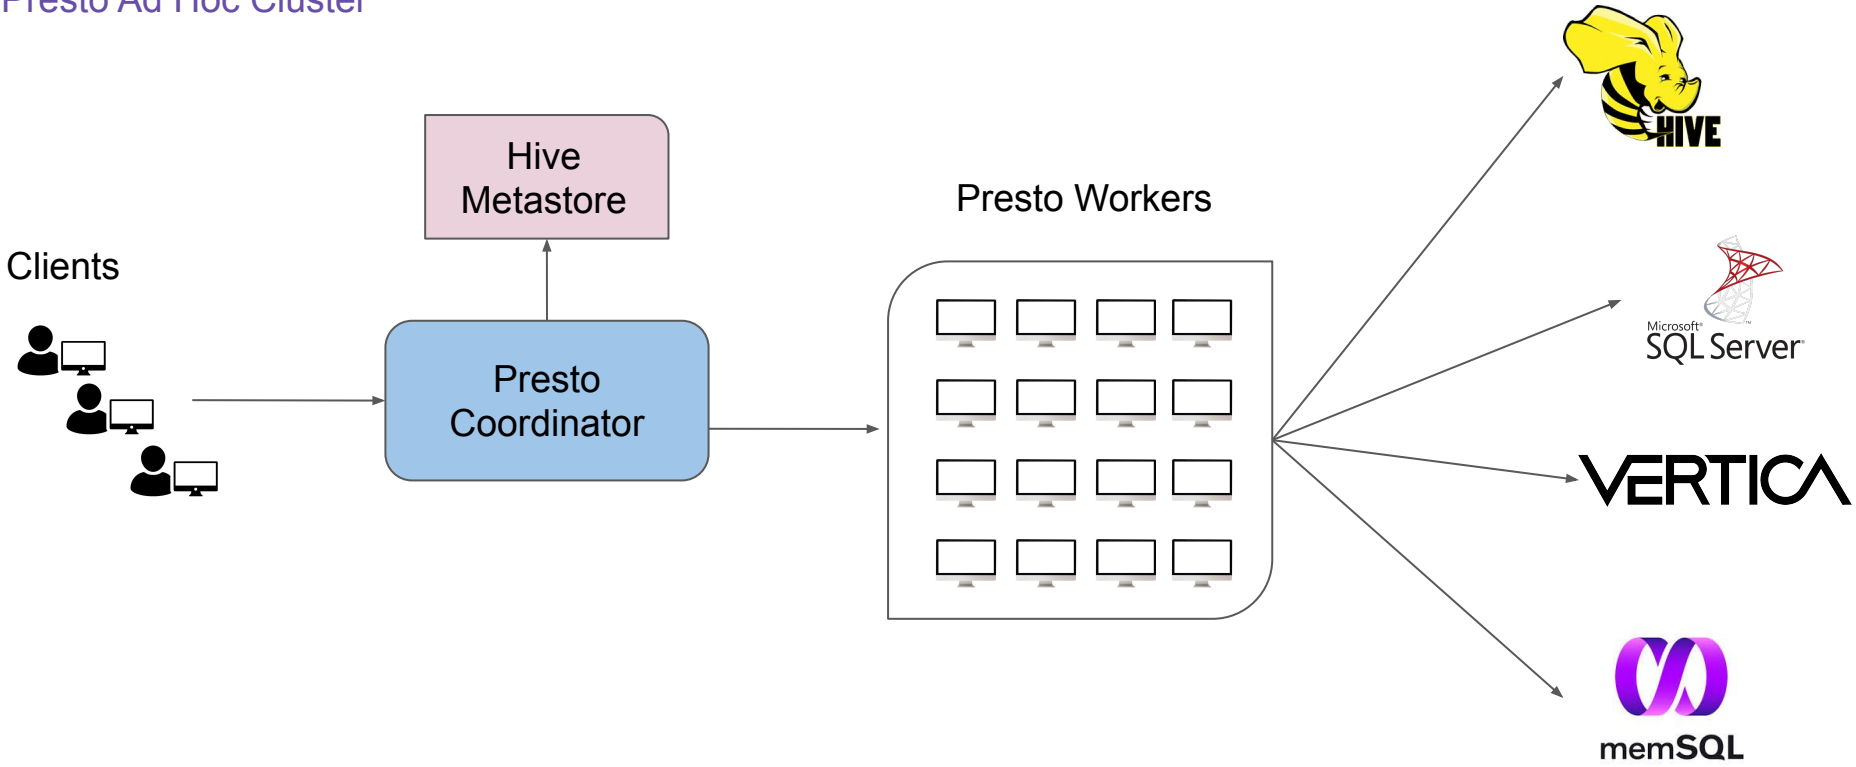

Presto At Wayfair

Presto Ad Hoc Cluster

Presto ad hoc (Read Only Cluster)

Version: 0.217

301 VM's (8*64) with 1 Coordinator, 300 Workers

Total available Memory ~20TB

Total CPU available 2400 vcores

Adoption – before & after

Presto Ad Hoc Cluster

Hive **Queries per Month** prior to Presto

Presto's **performance** won **almost half** of Hive activity in **just two months.**

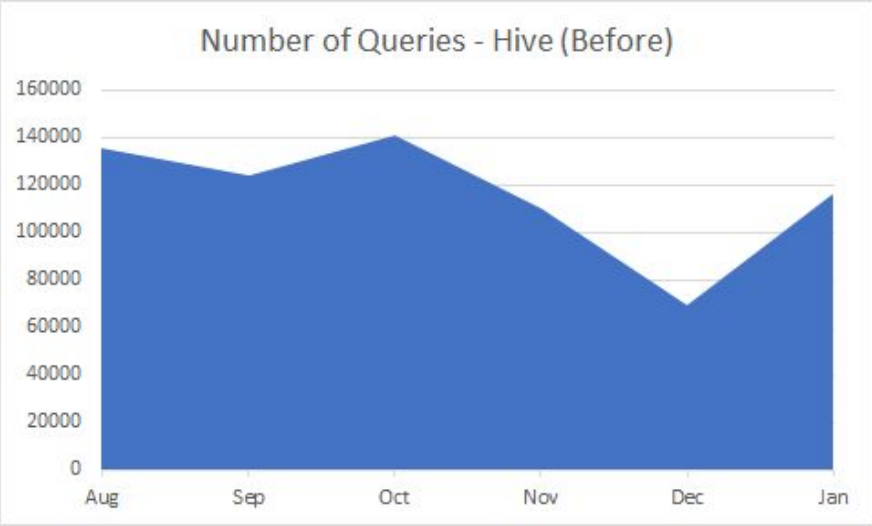

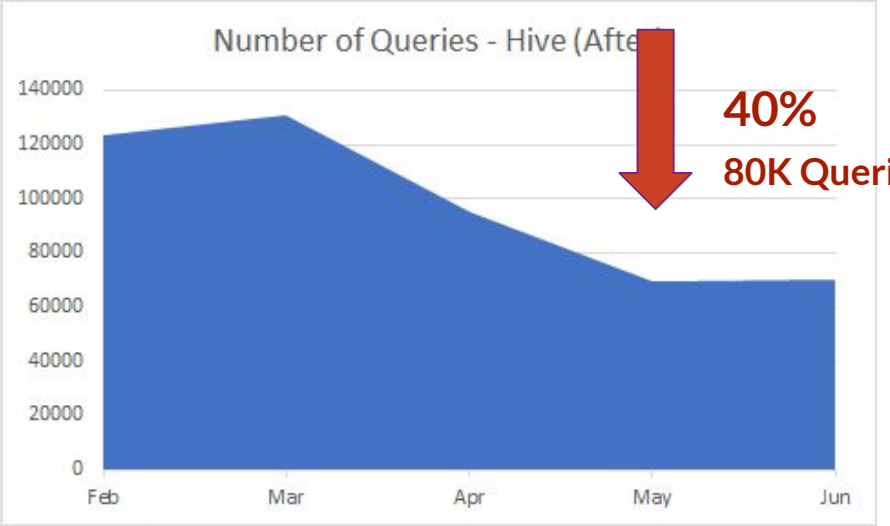

Adoption – after

Presto Ad Hoc Cluster

Presto **Queries per Month**

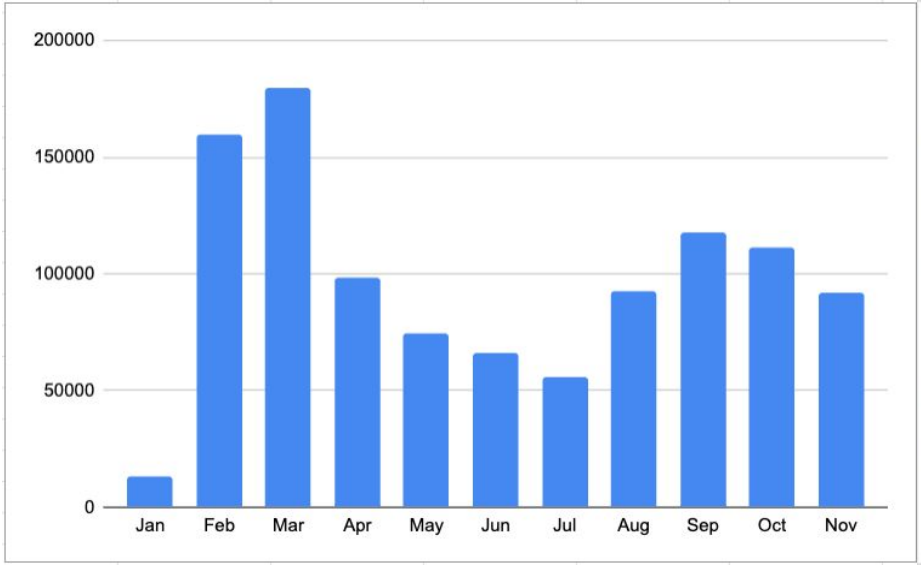

Presto **users growth** over the year

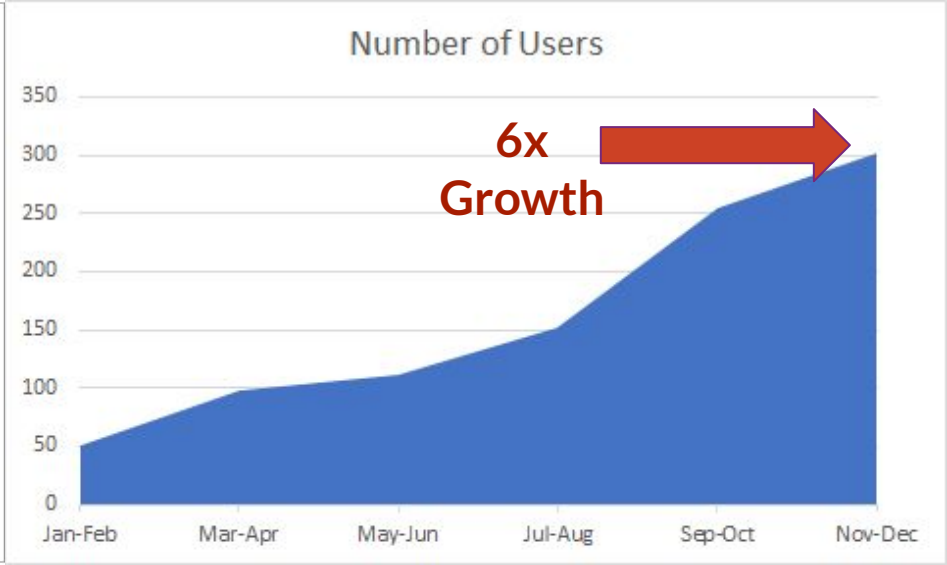

Query Throttling

Presto Ad Hoc Cluster

Year_Month	Mean_Runtime_seconds	Query_count
2019-01	11.839540	12906
2019-02	17.831149	160019
2019-03	14.425602	179657
2019-04	10.334873	98401
2019-05	48.924414	74685
2019-06	49.777620	65961
2019-07	51.911089	55706
2019-08	29.966837	71547
2019-09	13.857432	117773
2019-10	24.810308	111457
2019-11	15.742889	92030
2019-12	13.634272	15919

- SELECT only
- 2 queries per user
- 2 queued queries per user
- Increased the time limit from 5 to 10 mins

Avg execution time dropped from 51 secs to 20 secs

Presto Read/Write Cluster Beta

- In beta with 50 nodes
- Limited # of users using a default namespace
- Faster writes/inserts than hive
- Resource grouping is enabled via queues
- 10 min limit on query execution time

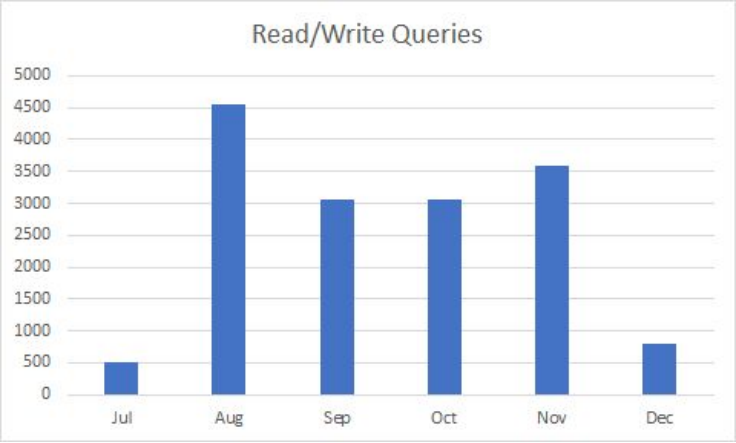

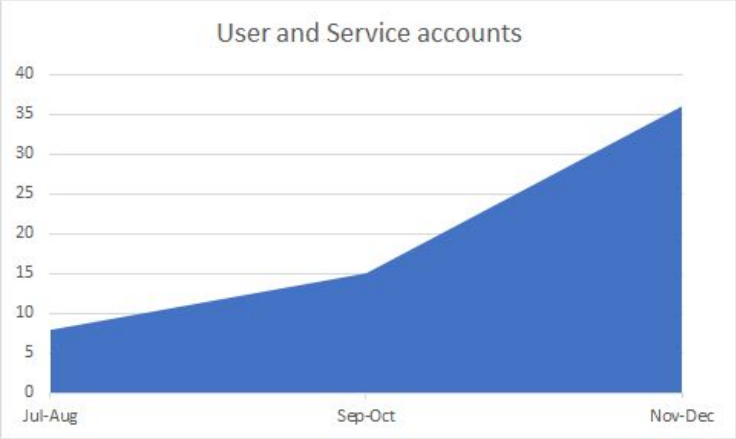

POC: Starburst Presto Distribution

presto → Starburst

Results

The results are summarized below. All times are in milliseconds; lower is better.

Cluster	version	Query 1			Query 2			Query 3			Query 4		
		Wall	Query	Process	Wall	Query	Process	Wall	Query	Process	Wall	Query	Process
APACHE PRESTO	0.217	764.81	26.02	738.55	153600	18.69	153600	5910	16.11	5900	86400	62.74	86400
STARBURST PRESTO	302-e.11	587.75	17.86	560.84	147000	18.11	147000	5760	15.62	5750	79200	72.5	79200

NOTE: Query5 exceeded the maximum execution time limit of 5.00m

Monitoring Presto

Skynet (internal)

What's Next

Continue migrating jobs to Presto

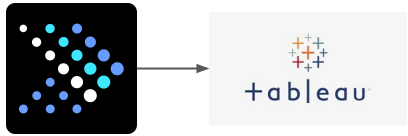

Presto to Tableau Connector

Presto in Google Cloud

Rationalize BigQuery vs Presto in GCP

Questions?

wayfair®

Adoption – surprises

Presto Ad Hoc Cluster

- Overall, very few surprises!
- No official Presto connector for Vertica (very popular at Wayfair) and Vertica internal libraries are closed... so we wrote our own connector
- Performance & unification has become so popular, devs now asking to point their applications to Presto as a data interface layer (evaluating…)