

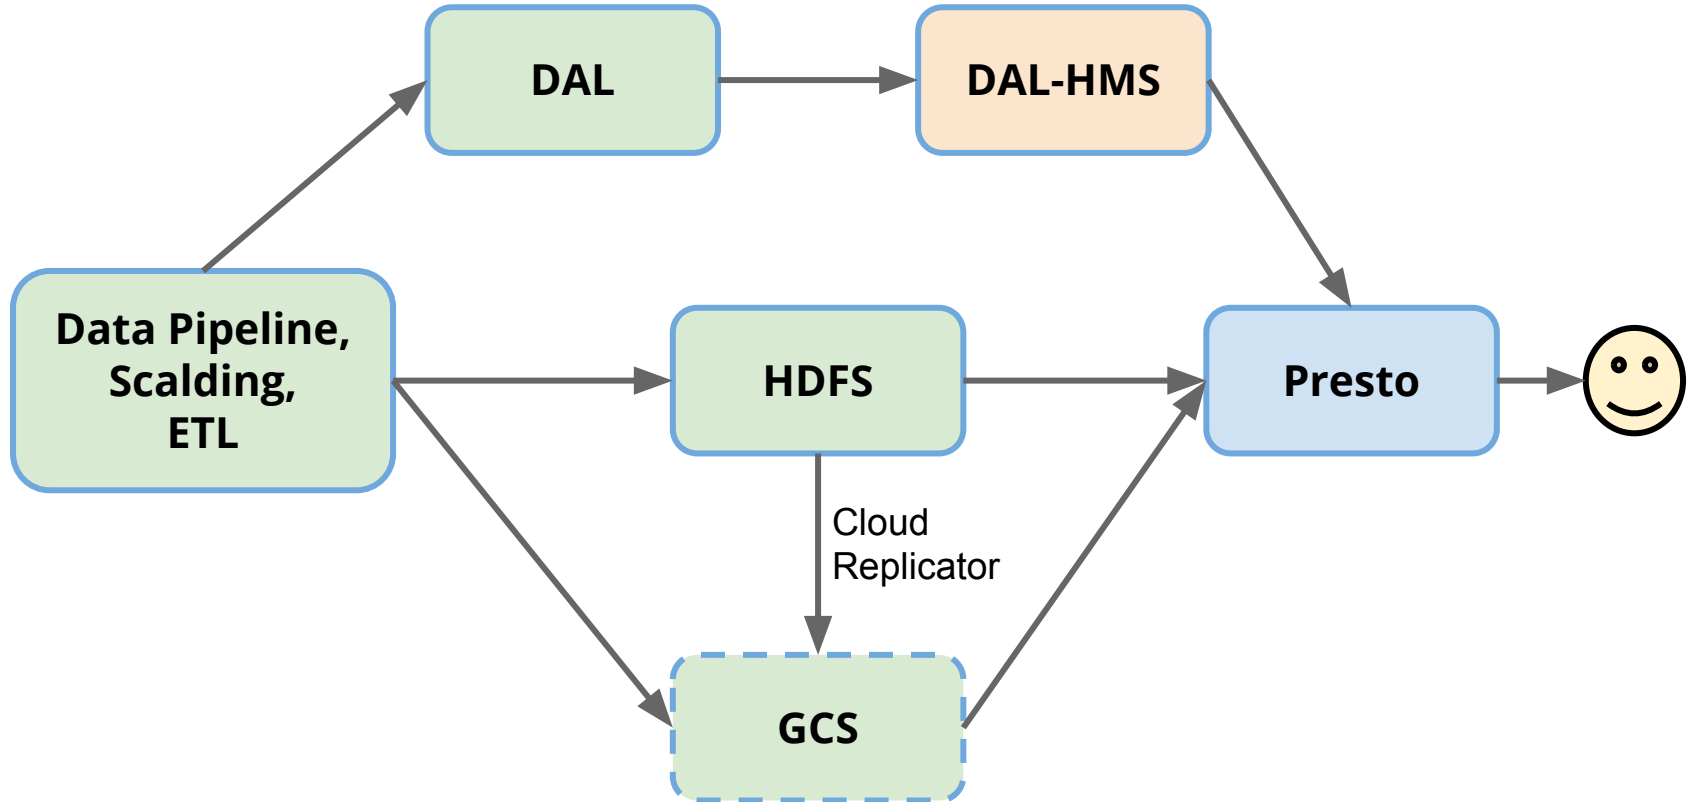
Presto@Twitter

Maosong Fu - mfu@twitter.com





Analytics Pipeline @Twitter





Presto @Twitter

- 4 Clusters on-prem
 - 200 ~ 500 nodes each
 - deployed in Aurora
- 1 Cluster in the cloud
 - up to 800 nodes
 - deployed in DataProc



Presto Query @Twitter

- Ad-hoc query only
- Data format: Parquet, lzo-thrift
- Daily queries: ~20K
- Daily processed data: ~30PB

Authentication, Authorization, Auditing and More





Authentication

- Kerberos
- LDAP



Authorization

- Storage-based security model
- HDFS
 - Hadoop Impersonation
 - Manage access via LDAP groups per cluster that presto can proxy as
- Google Cloud Storage
 - per-user authorization w/ OAuth token

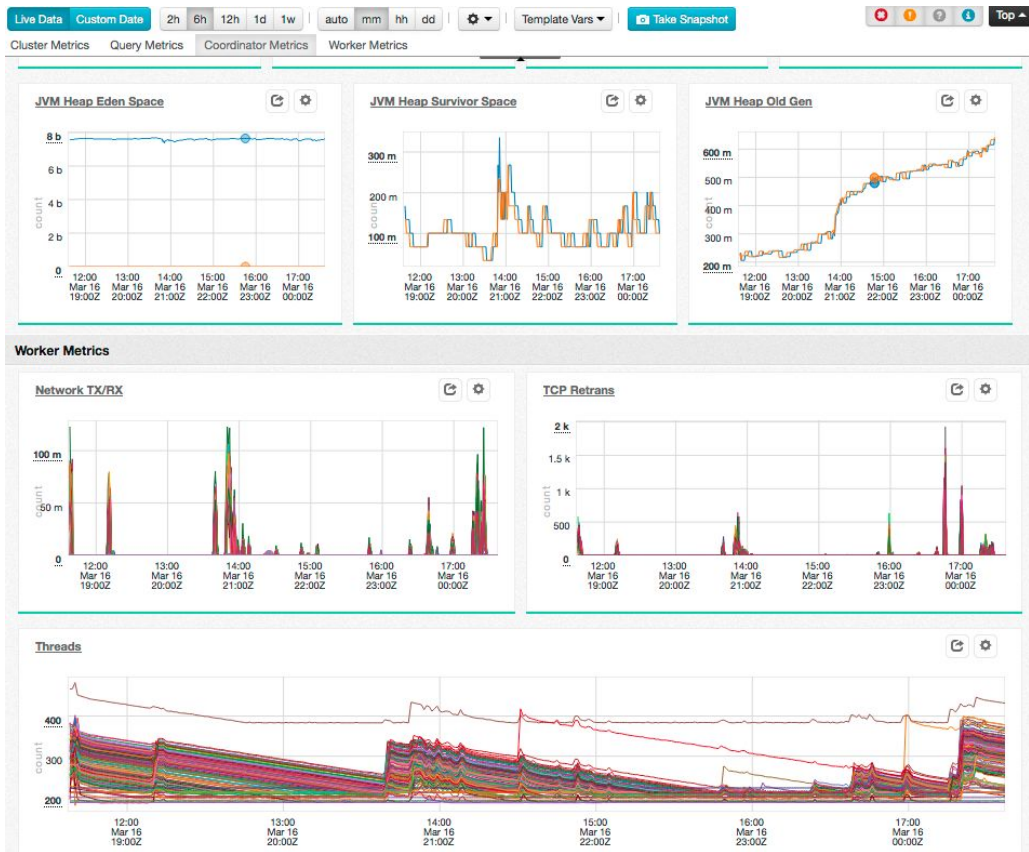


Auditing

- Query Completion Log
 - All query completion events go to HDFS log category
 - The Presto query completion log can be queried using Presto



Monitoring & Alerting



Stability





Live humongous region 623 size 43217136 ...

- Memory fragmentation in JVM OG causing frequent OOMs
- Dynamic batch sizing in Parquet reader
 - Estimate the size of each cell of a the table based on history
 - Predict the batch size based on max combined cell size

Future Work





Presto Roadmap

- Cloud migration
- More connectors
 - MySQL, Vertica, BigQuery
 - pushdown of complex operations

Thank you.

Q&A

