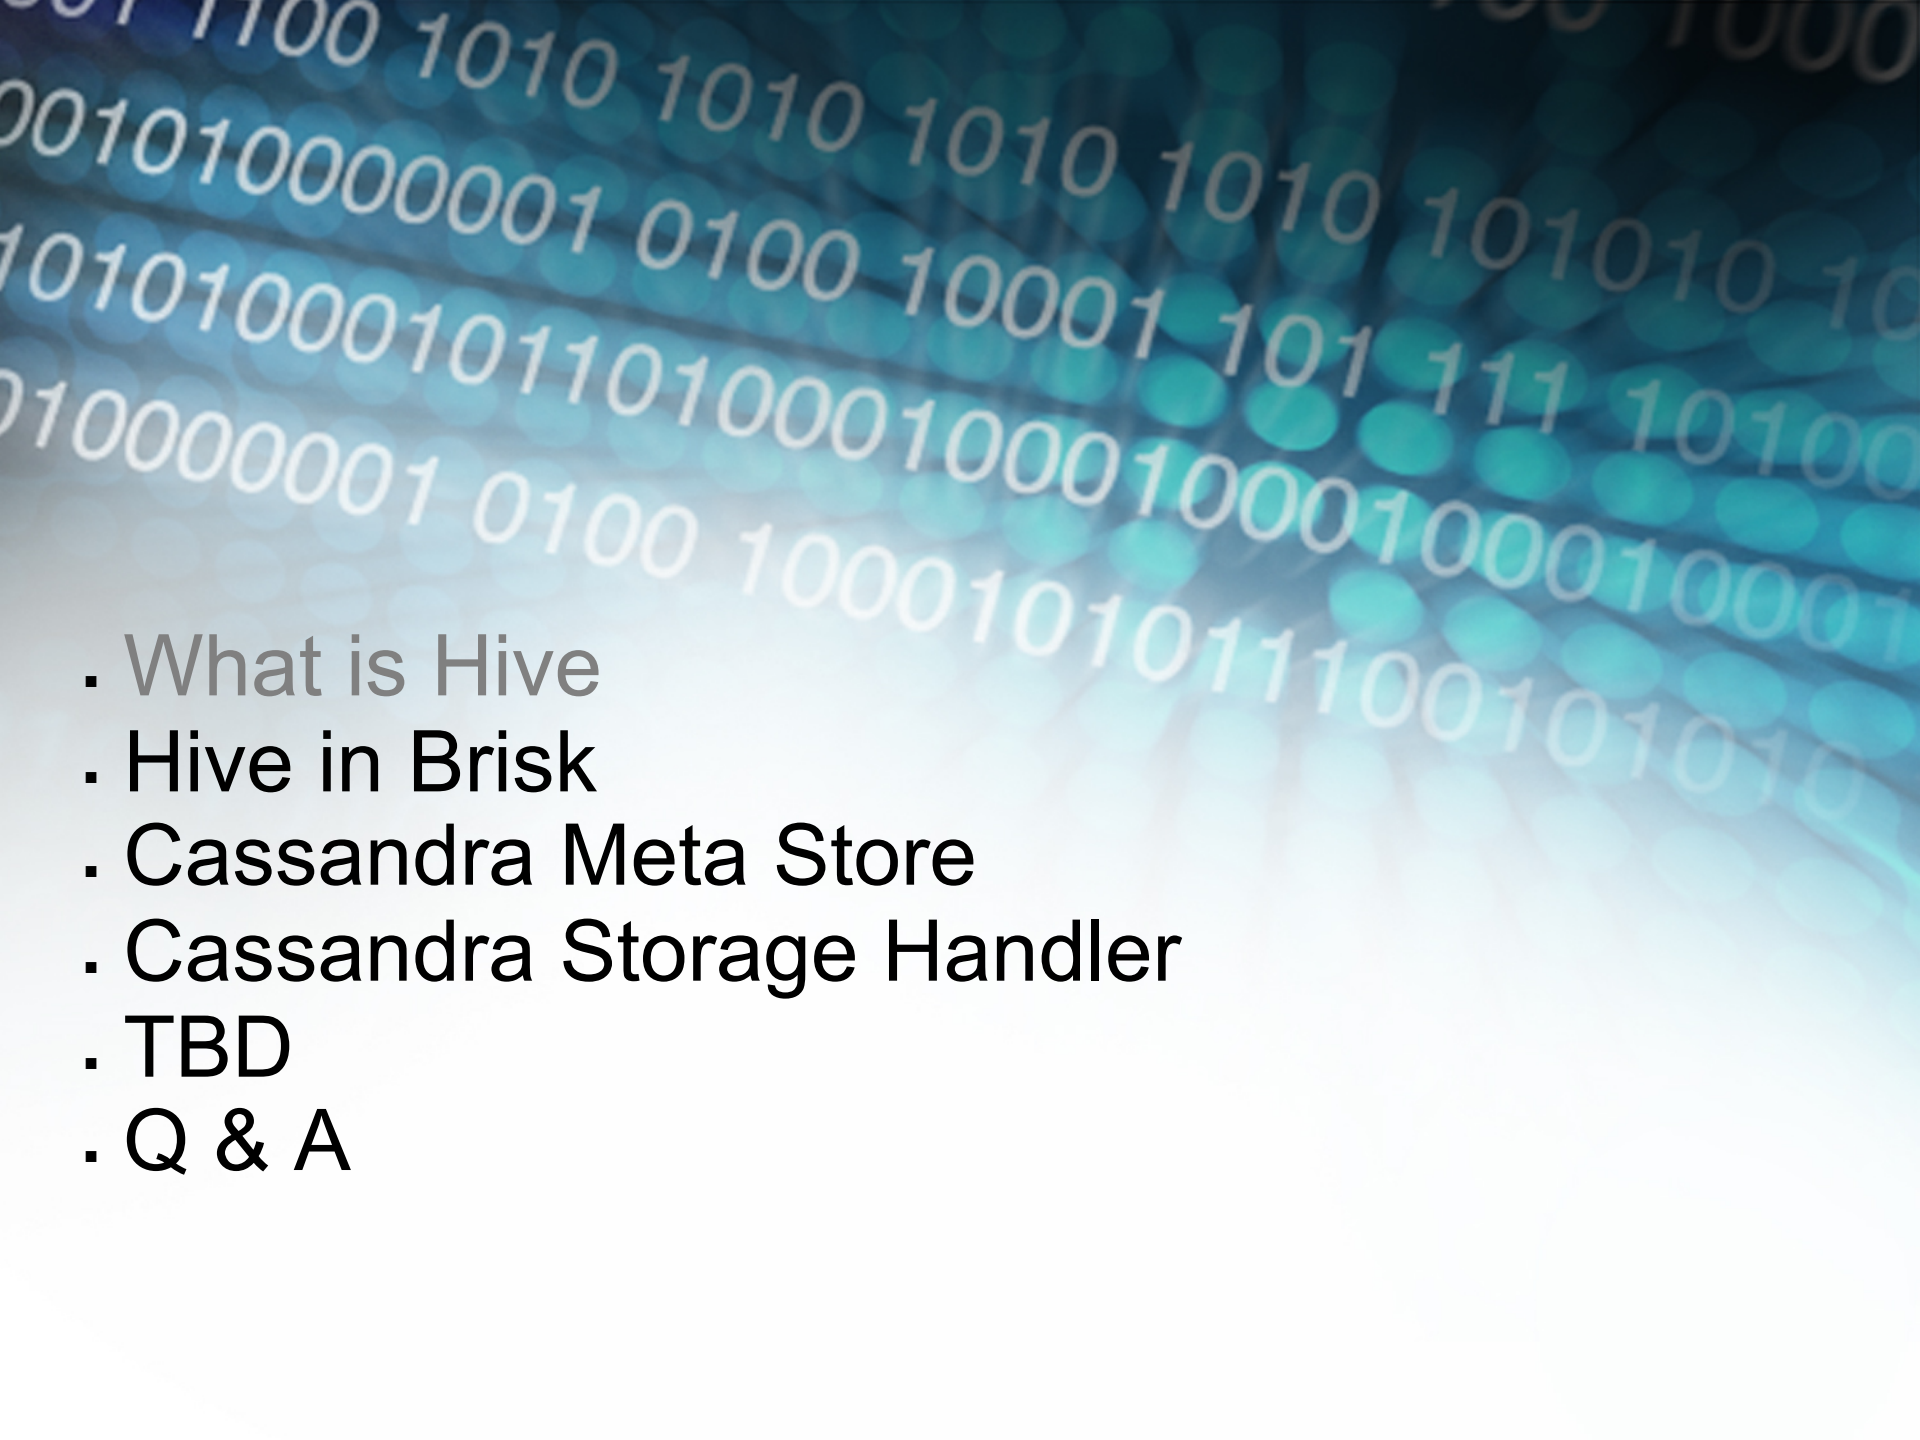


Cassandra SF 2011

Hive over Brisk

Yewei Zhang
Software Engineer

- 
- . What is Hive
 - . Hive in Brisk
 - . Cassandra Meta Store
 - . Cassandra Storage Handler
 - . TBD
 - . Q & A

What is Apache Hive?

Data warehouse infrastructure on top of Hadoop

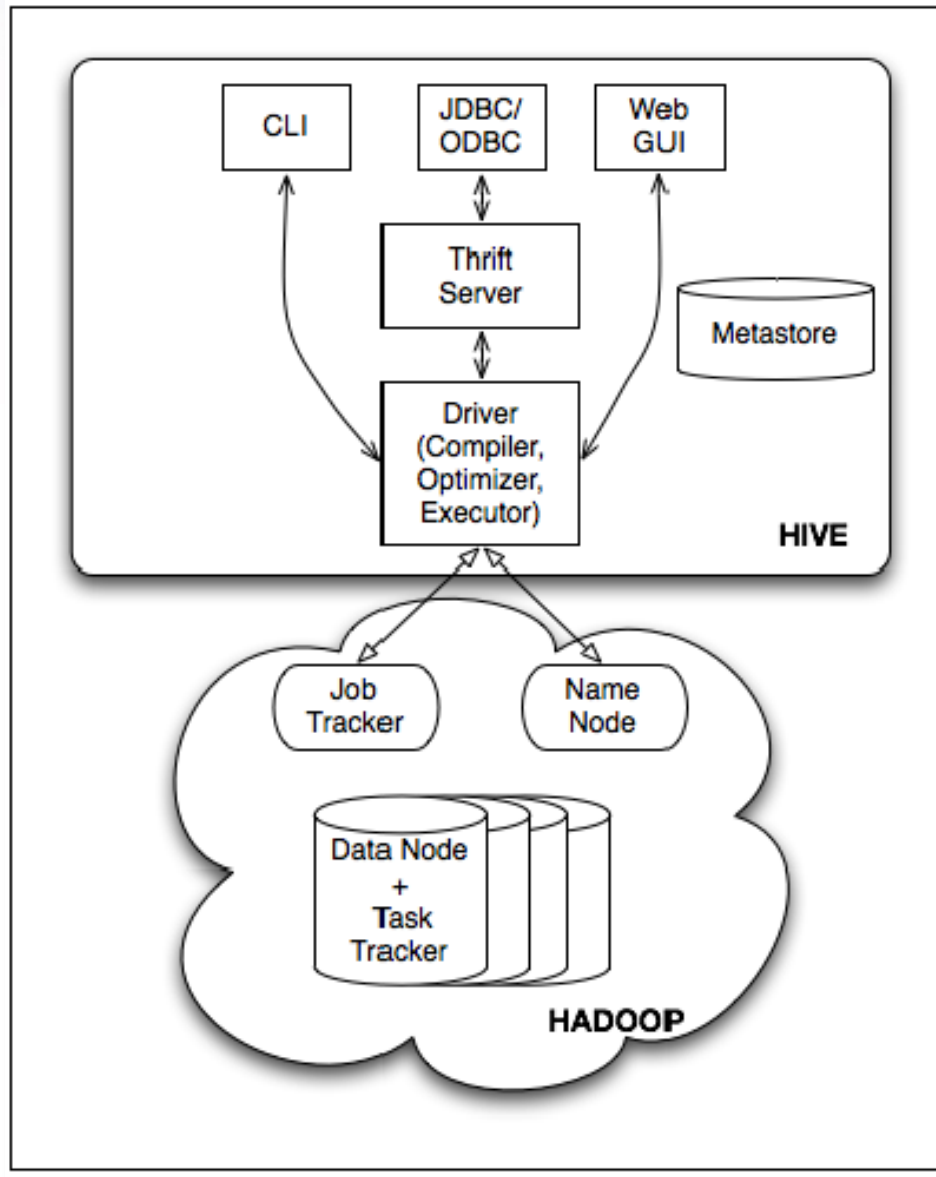
- .SQL-like query

- .Customized mapper and reducer plugin

- .Open data formats




Hive Architecture



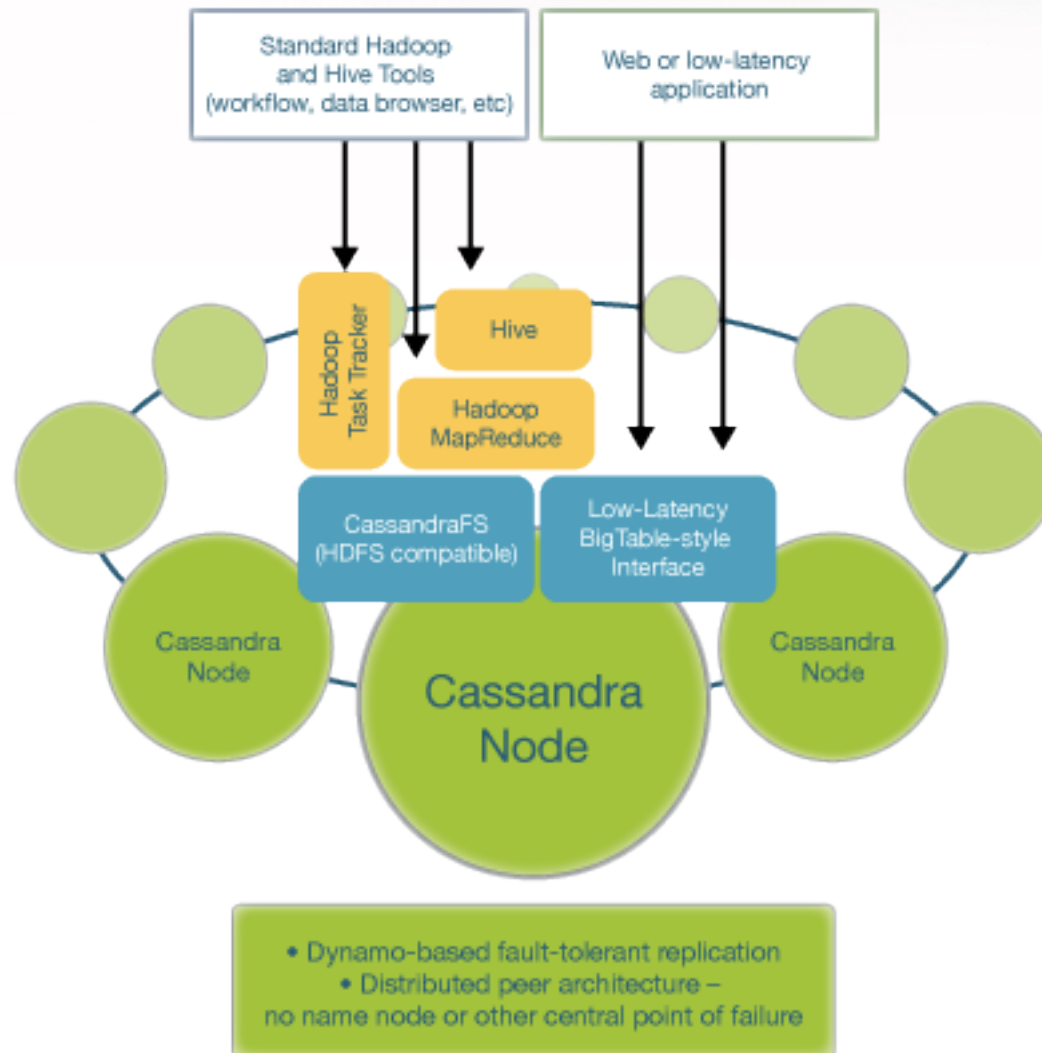
When to use Hive

- Batch jobs over large sets of immutable data
- Large Hadoop cluster
- Data has a schema

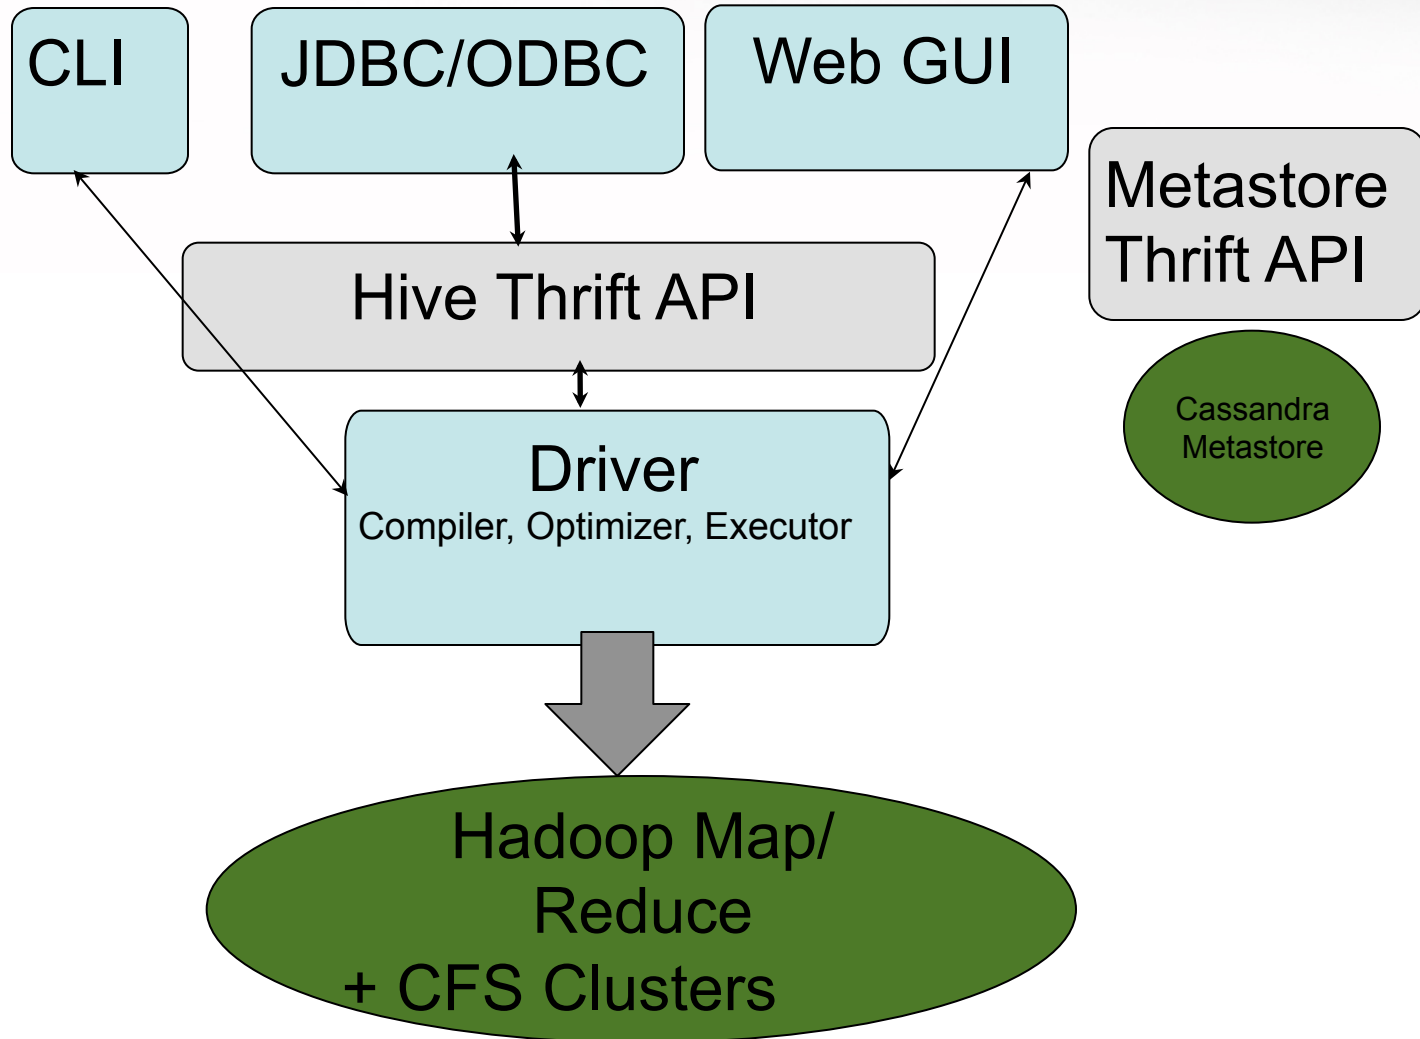



- 
- . What is Hive
 - . Hive in Brisk
 - . Cassandra Meta Store
 - . Cassandra Storage Handler
 - . TBD
 - . Q & A

Dramatically Simplified Hadoop and Low-Latency Application Stack

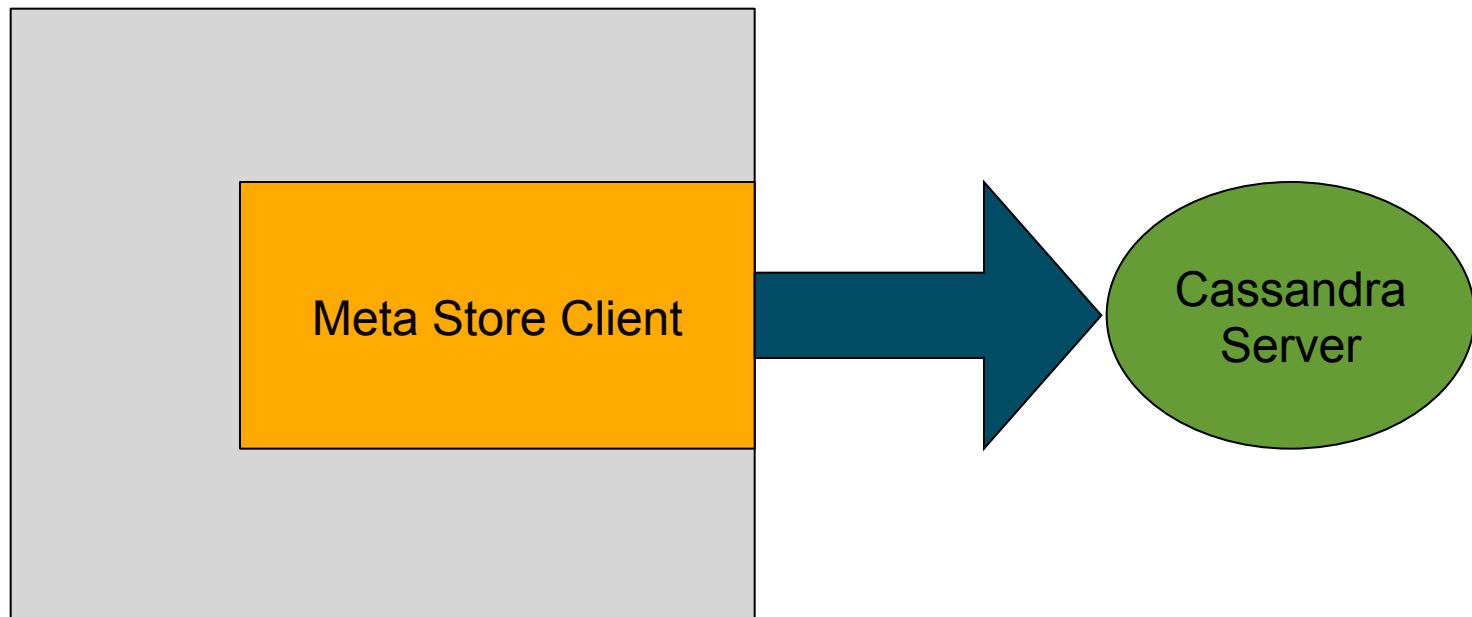


Hive in Brisk



- 
- . What is Hive
 - . Hive in Brisk
 - . Cassandra Meta Store
 - . Cassandra Storage Handler
 - . TBD
 - . Q & A

Cassandra MetaStore



Cassandra MetaStore

hive-site.xml

```
<property>
  <name>hive.metastore.warehouse.dir</name>
  <value>cfs:///user/hive/warehouse</value>
</property>
```

...

```
<property>
  <name>hive.metastore.rawstore.impl</name>
  <value>
org.apache.cassandra.hadoop.hive.metastore.CassandraHiveMeta
Store
  </value>
</property>
```

Cassandra MetaStore Configuration

- `cassandra.connection.host`
- `cassandra.connection.port`
- `cassandra.connection.metaStoreKeyspaceName`
- `cassandra.connection.readCL`
- `cassandra.connection.writeCL`
- `cassandra.connection.replicationFactor`

Cassandra MetaStore Configuration

hive-site.xml

```
<property>  
<name>cassandra.connection.host</name>  
<value>[your brisk node host]</value>  
</property>
```

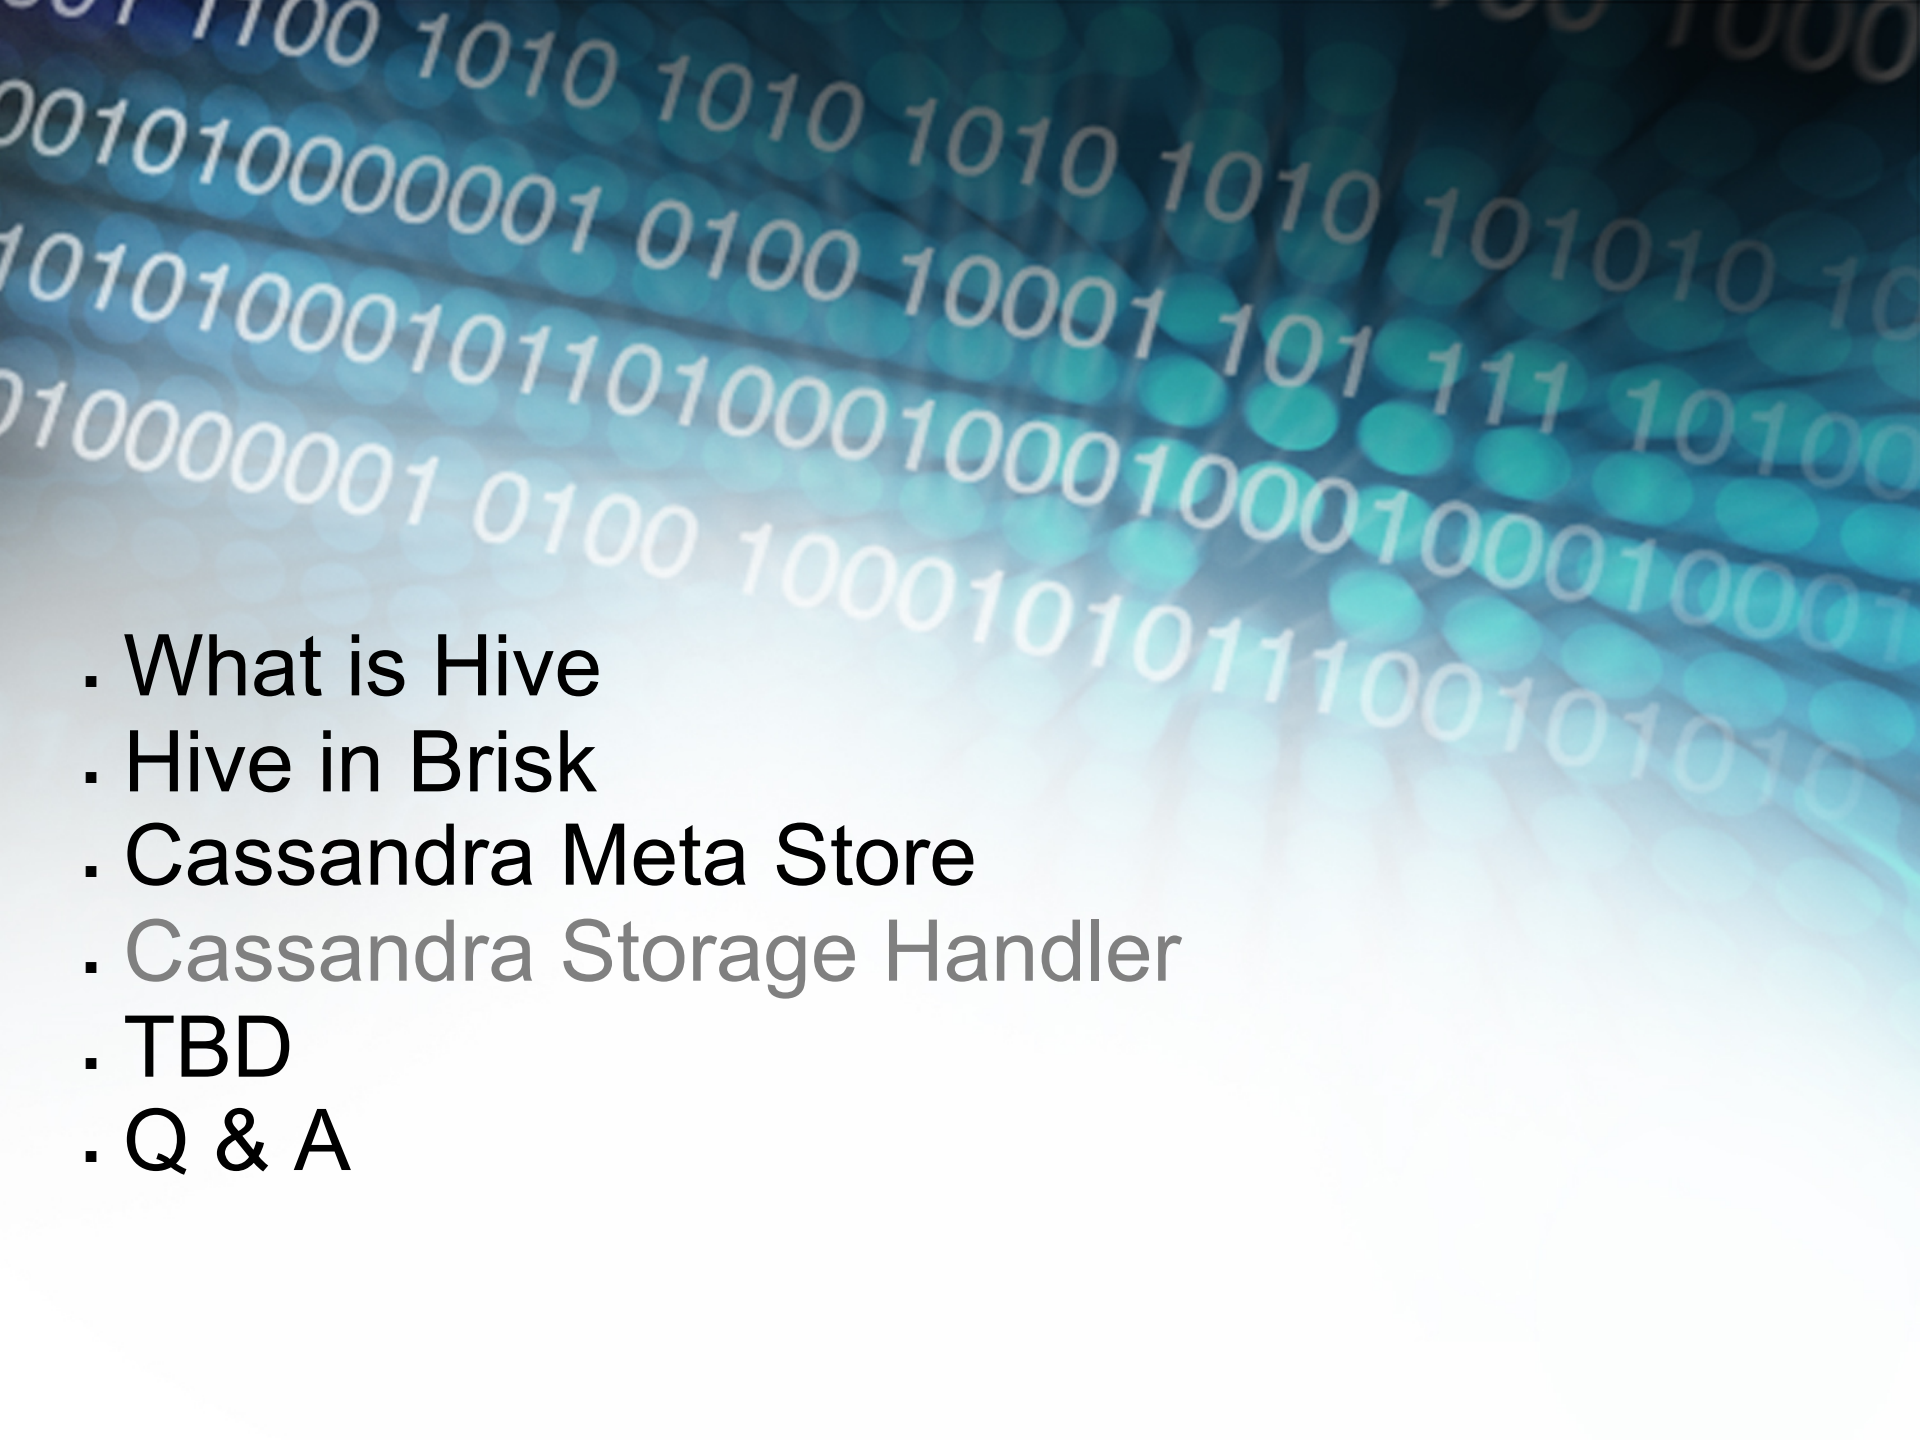
```
<property>  
<name>cassandra.connection.port</name>  
<value>[your brisk node port]</value>  
</property>
```

...

External tables auto mapping

hive-site.xml

```
<!-- Set this to true to enable auto-creation of Cassandra  
keyspaces as Hive Databases -->  
<property>  
  <name>cassandra.autoCreateHiveSchema</name>  
  <value>true</value>  
</property>
```

- 
- . What is Hive
 - . Hive in Brisk
 - . Cassandra Meta Store
 - . Cassandra Storage Handler
 - . TBD
 - . Q & A

Cassandra Storage Handler

```
CREATE EXTERNAL TABLE
cassandra_hive_table(key int, column string, value string)

STORED BY
'org.apache.hadoop.hive.cassandra.CassandraStorageHandler'

WITH SERDEPROPERTIES (
"cassandra.columns.mapping" = ":key,column,value" ,
"cassandra.cf.name" = "Table" ,
"cassandra.host" = "127.0.0.1" ,
"cassandra.port" = "9170")

TBLPROPERTIES (
"cassandra.ks.name" = "Hive");
```

Transposed table - column family

```
CREATE EXTERNAL TABLE divide  
  (row_key string,  
   column_name string,  
   value string)
```

```
STORED BY
```

```
'org.apache.hadoop.hive.cassandra.CassandraStorageHandler'
```

```
WITH SERDEPROPERTIES ("cassandra.host" = "127.0.0.1" ,  
"cassandra.port" = "9170")
```

```
TBLPROPERTIES ("cassandra.ks.name" = "Hive");
```

Transposed table - column family

row_key	column name	value
21	divide7	3
21	divide3	7

RowKey: 21

=> (column=divide7, value=3, timestamp=1307138105311000)

=> (column=divide3, value=7, timestamp=1307138105311000)

Transposed table - Super CF

```
CREATE EXTERNAL TABLE divide
    (key string, operator string, factor int, value string)
    STORED BY
    'org.apache.hadoop.hive.cassandra.CassandraStorageHandler'

WITH SERDEPROPERTIES (
    "cassandra.host" = "127.0.0.1" ,
    "cassandra.port" = "9170",
    "cassandra.columns.mapping" =
        ":key, :column, :subcolumn, :value")

TBLPROPERTIES ("cassandra.ks.name" = "Hive");
```

Transposed table - Super CF

row_key	operator	factor	value
21	divide	3	7
21	divide	7	3

RowKey: 21

=> (super_column=divide,
(column=7, value=3, timestamp=1307138105311000)
(column=3, value=7, timestamp=1307138105311000))

Transposed Table

cassandra.column.mapping

.:key

.:column

.:subcolumn

.:value

Transposed Table

Column Names

.row_key

.column_name

.sub_column_name

.value

External tables auto mapping

Mapped to transposed tables

ColumnFamily: divide

RowKey: 21

=> (column=divide7, value=3, timestamp=1307138105311000)

=> (column=divide3, value=7, timestamp=1307138105311000)

External tables auto mapping

Mapped to transposed tables

```
hive> desc event_activity_day
row_key          string          from deserializer
column_name      string          from deserializer
value           string          from deserializer
```

Configuration Properties

· SERDEPROPERTIES

- cassandra.columns.mapping
- cassandra.cf.name
- cassandra.host
- cassandra.port

· TBLPROPERTIES

- cassandra.ks.name
- cassandra.range.size
- cassandra.input.split.size
- ...

Column Mapping

`cassandra.columns.mapping`

`.:key` - column family row key

`.:column` - column names

`.:subcolumn` - subcolumn names (in super CF)

`.:value` - column/subcolumn values

Transposed Mapping

```
hive> CREATE EXTERNAL TABLE
      PortfolioDemo.PortfolioStocks(
          portfolio string,
          ticker string,
          number_shares string)
      STORED BY
      'org.apache.hadoop.hive.cassandra.CassandraStorageHandler'

      WITH SERDEPROPERTIES (
          "cassandra.columns.mapping" =
          ":key, :column, :value" );
```


General Mapping

```
hive> CREATE EXTERNAL TABLE Users(  
    userid string,  
    name string,  
    email string,  
    phone string)  
STORED BY ...  
WITH SERDEPROPERTIES (  
    "cassandra.columns.mapping" = "  
        :key,  
        user_name,  
        primary_email,  
        home_phone" );
```

No Column Mapping

```
hive> CREATE EXTERNAL TABLE Users(  
    userid string,  
    name string,  
    email string,  
    phone string)  
    STORED BY ...  
    WITH SERDEPROPERTIES (...);
```

- . **1st table field** - column family row key

- 
- . What is Hive
 - . Hive in Brisk
 - . Cassandra Meta Store
 - . Cassandra Storage Handler
 - . TBD
 - . Q & A

Future Works

- Pushdown predicate support
- User/Group/Role & Privileges support
- Performance comparison between Hive and Hive in Brisk

Pushdown Predicate Support


Select * from <<Table>>

where

x > 3 AND upper(y) = 'XYZ'

Pushdown Predicate Support

- <https://cwiki.apache.org/confluence/display/Hive/FilterPushdownDev>
- Push filters down to storage handler
- Take advantage of the secondary indexing in Cassandra
- Depend on Cassandra-1600

- 
- What is Hive
 - Hive in Brisk
 - Cassandra Meta Store
 - Cassandra Storage Handler
 - TBD
 - Q & A

References

- DataStax Brisk wiki page

http://www.datastax.com/docs/0.8/brisk/about_hive

- Hive wiki page

<https://cwiki.apache.org/confluence/display/Hive/Home>

Resources

- <http://www.datastax.com/docs/0.8/brisk/index>
- <https://github.com/riptano/hive>
- brisk-users@googlegroups.com
- freenode #datastax-brisk
- yewei@datastax.com