

*Inspire
Creativity
Enrich
Life*

 ByteDance 字节跳动

Presto@ByteDance

Pengfei Chang
2021 Dec.

01

Presto at Bytedance OLAP platform

02

Presto Cluster Stability and Maintainability Improvement

03



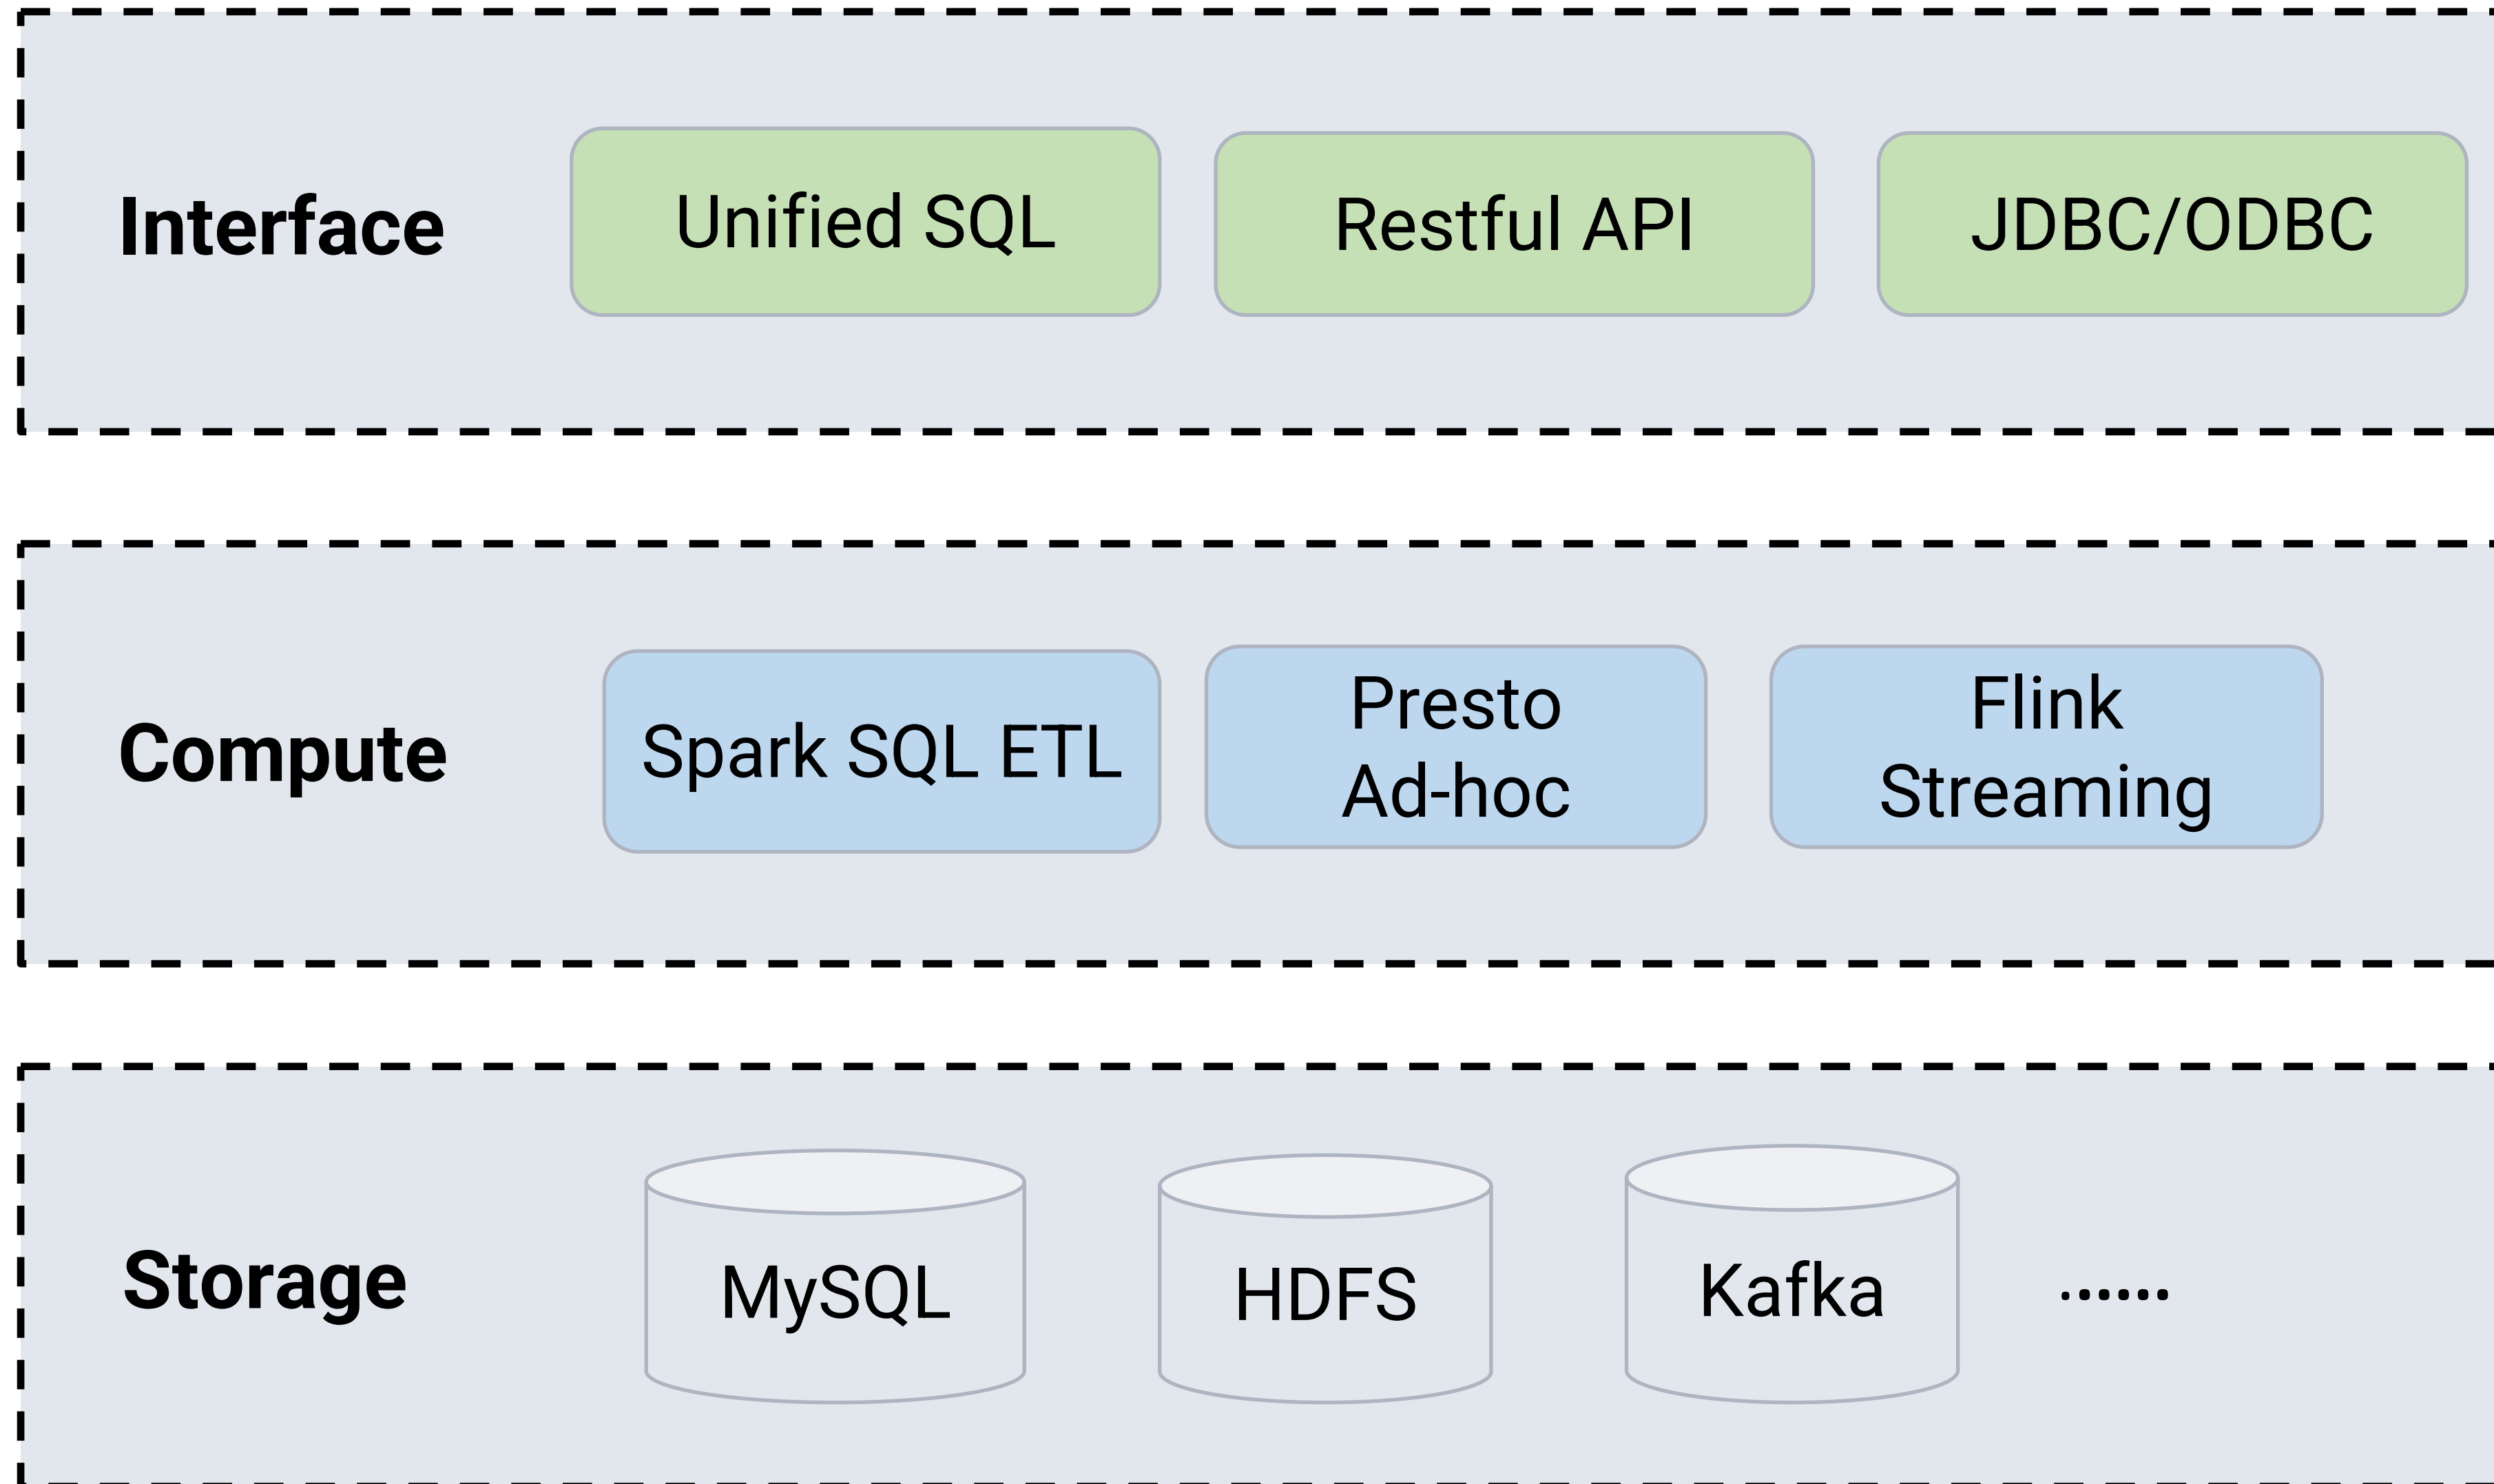
Practices and Optimizations at Different Scenarios

04

Ongoing Exploration and Future Plan


1

Presto at Bytedance OLAP Platform

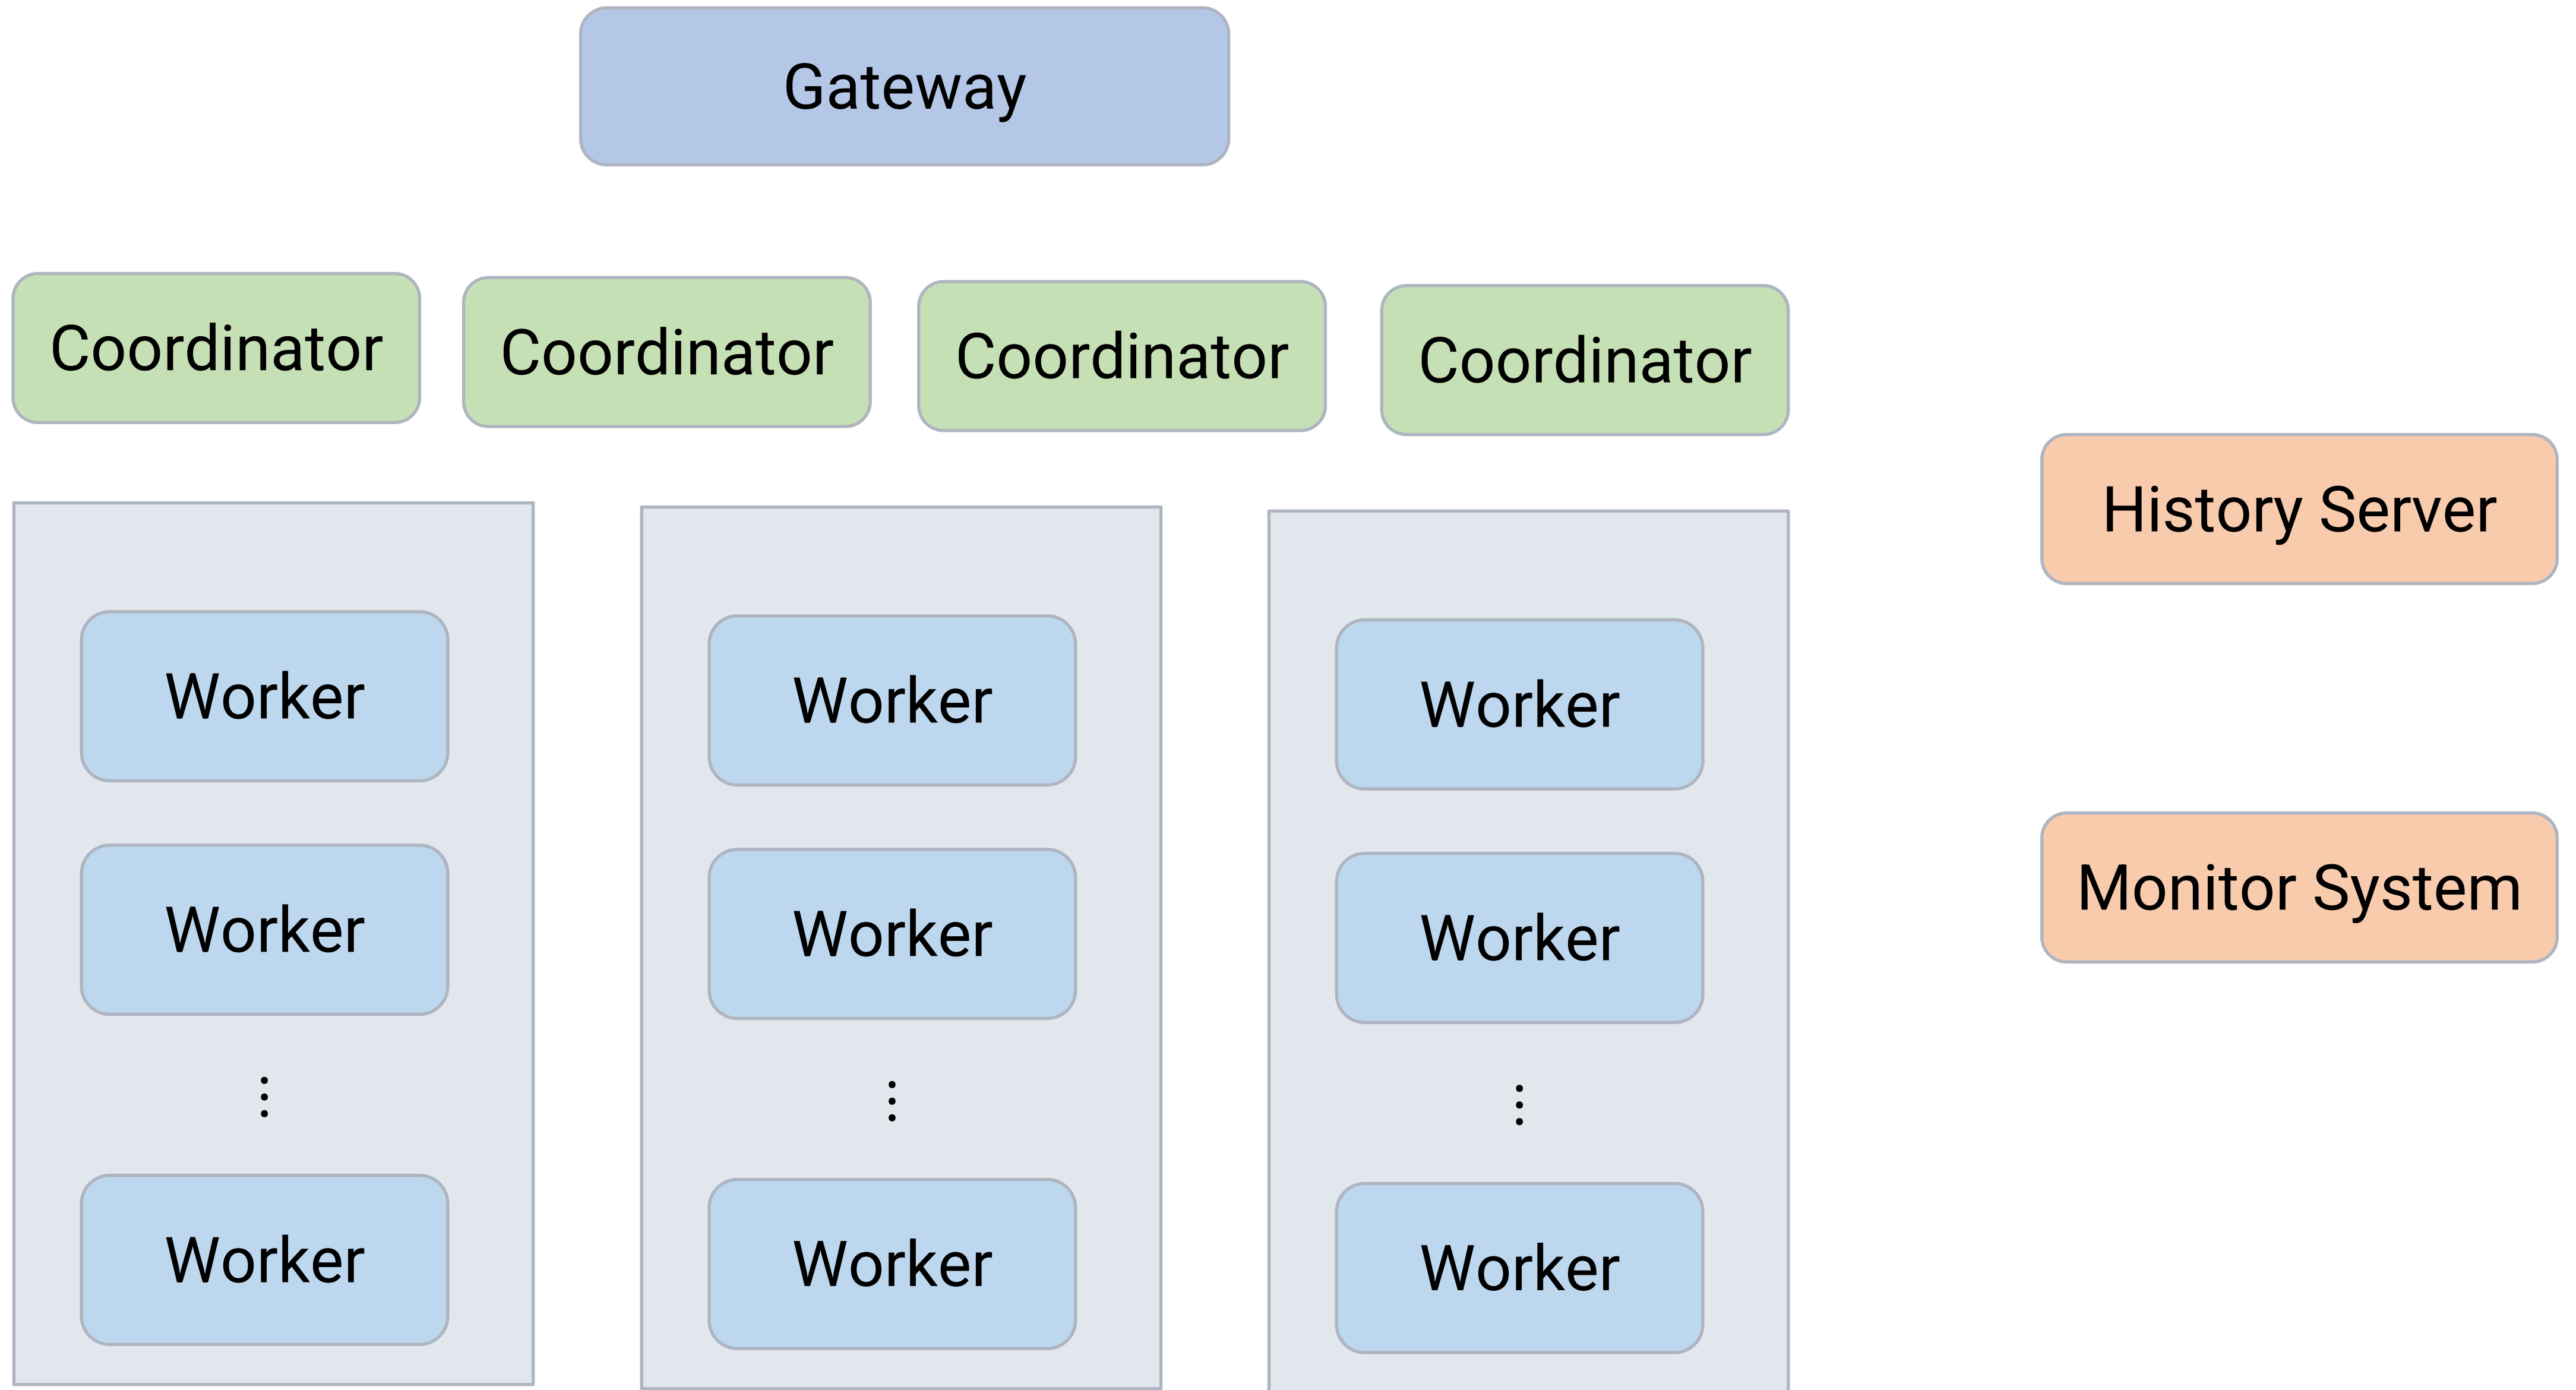


Presto at Bytedance

- Started from 2020
- Tens of thousands compute cores
- ~ 1 million queries per day
- Cover BI, Ad-hoc scenarios



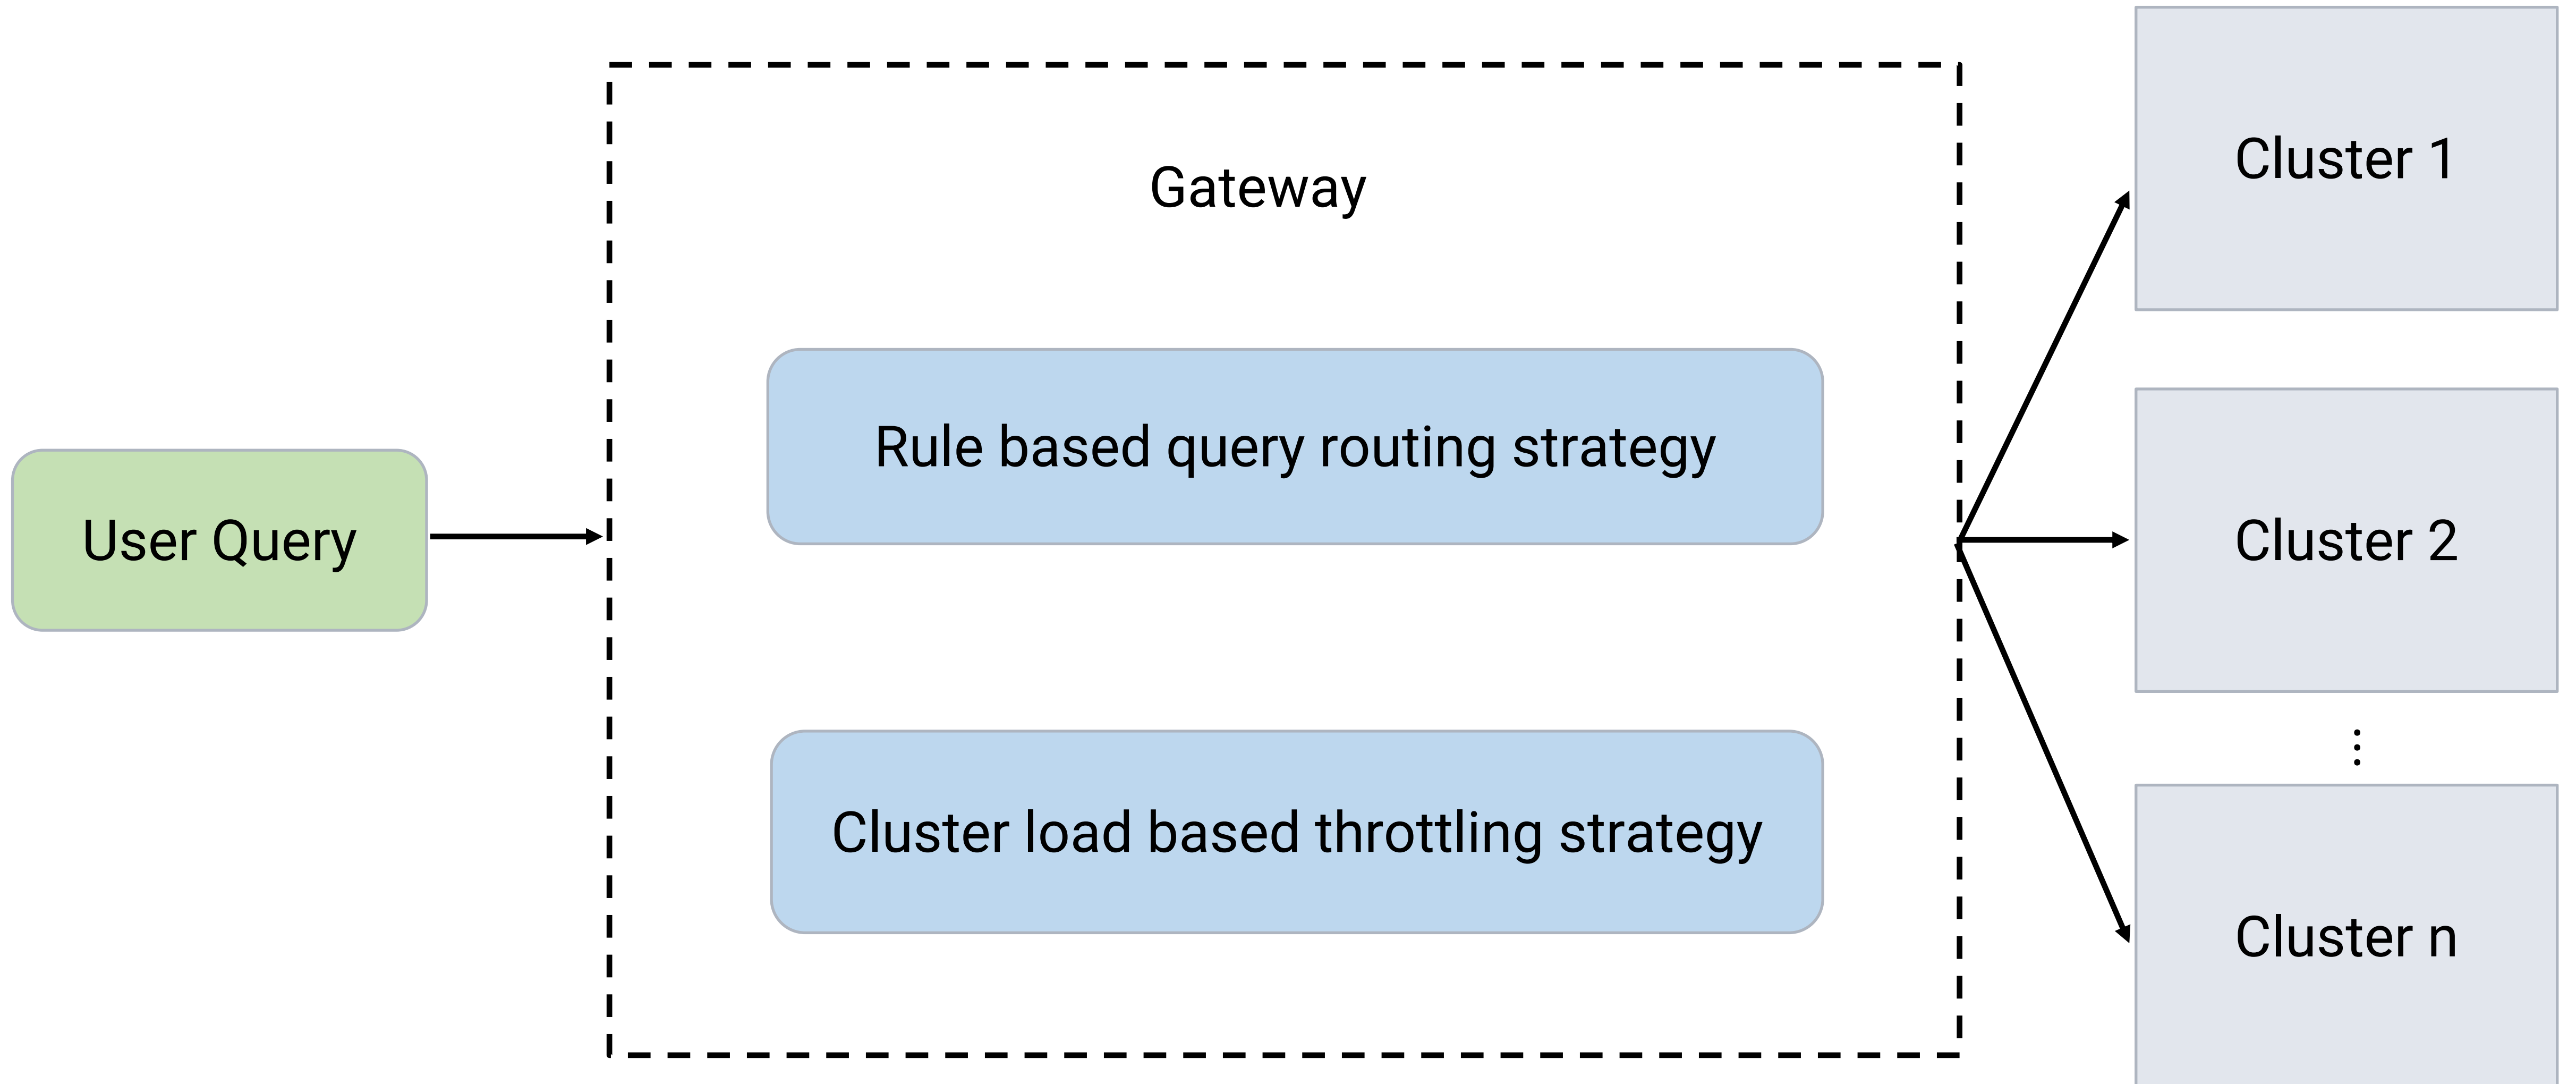
Presto Queries Per Day Since 2020



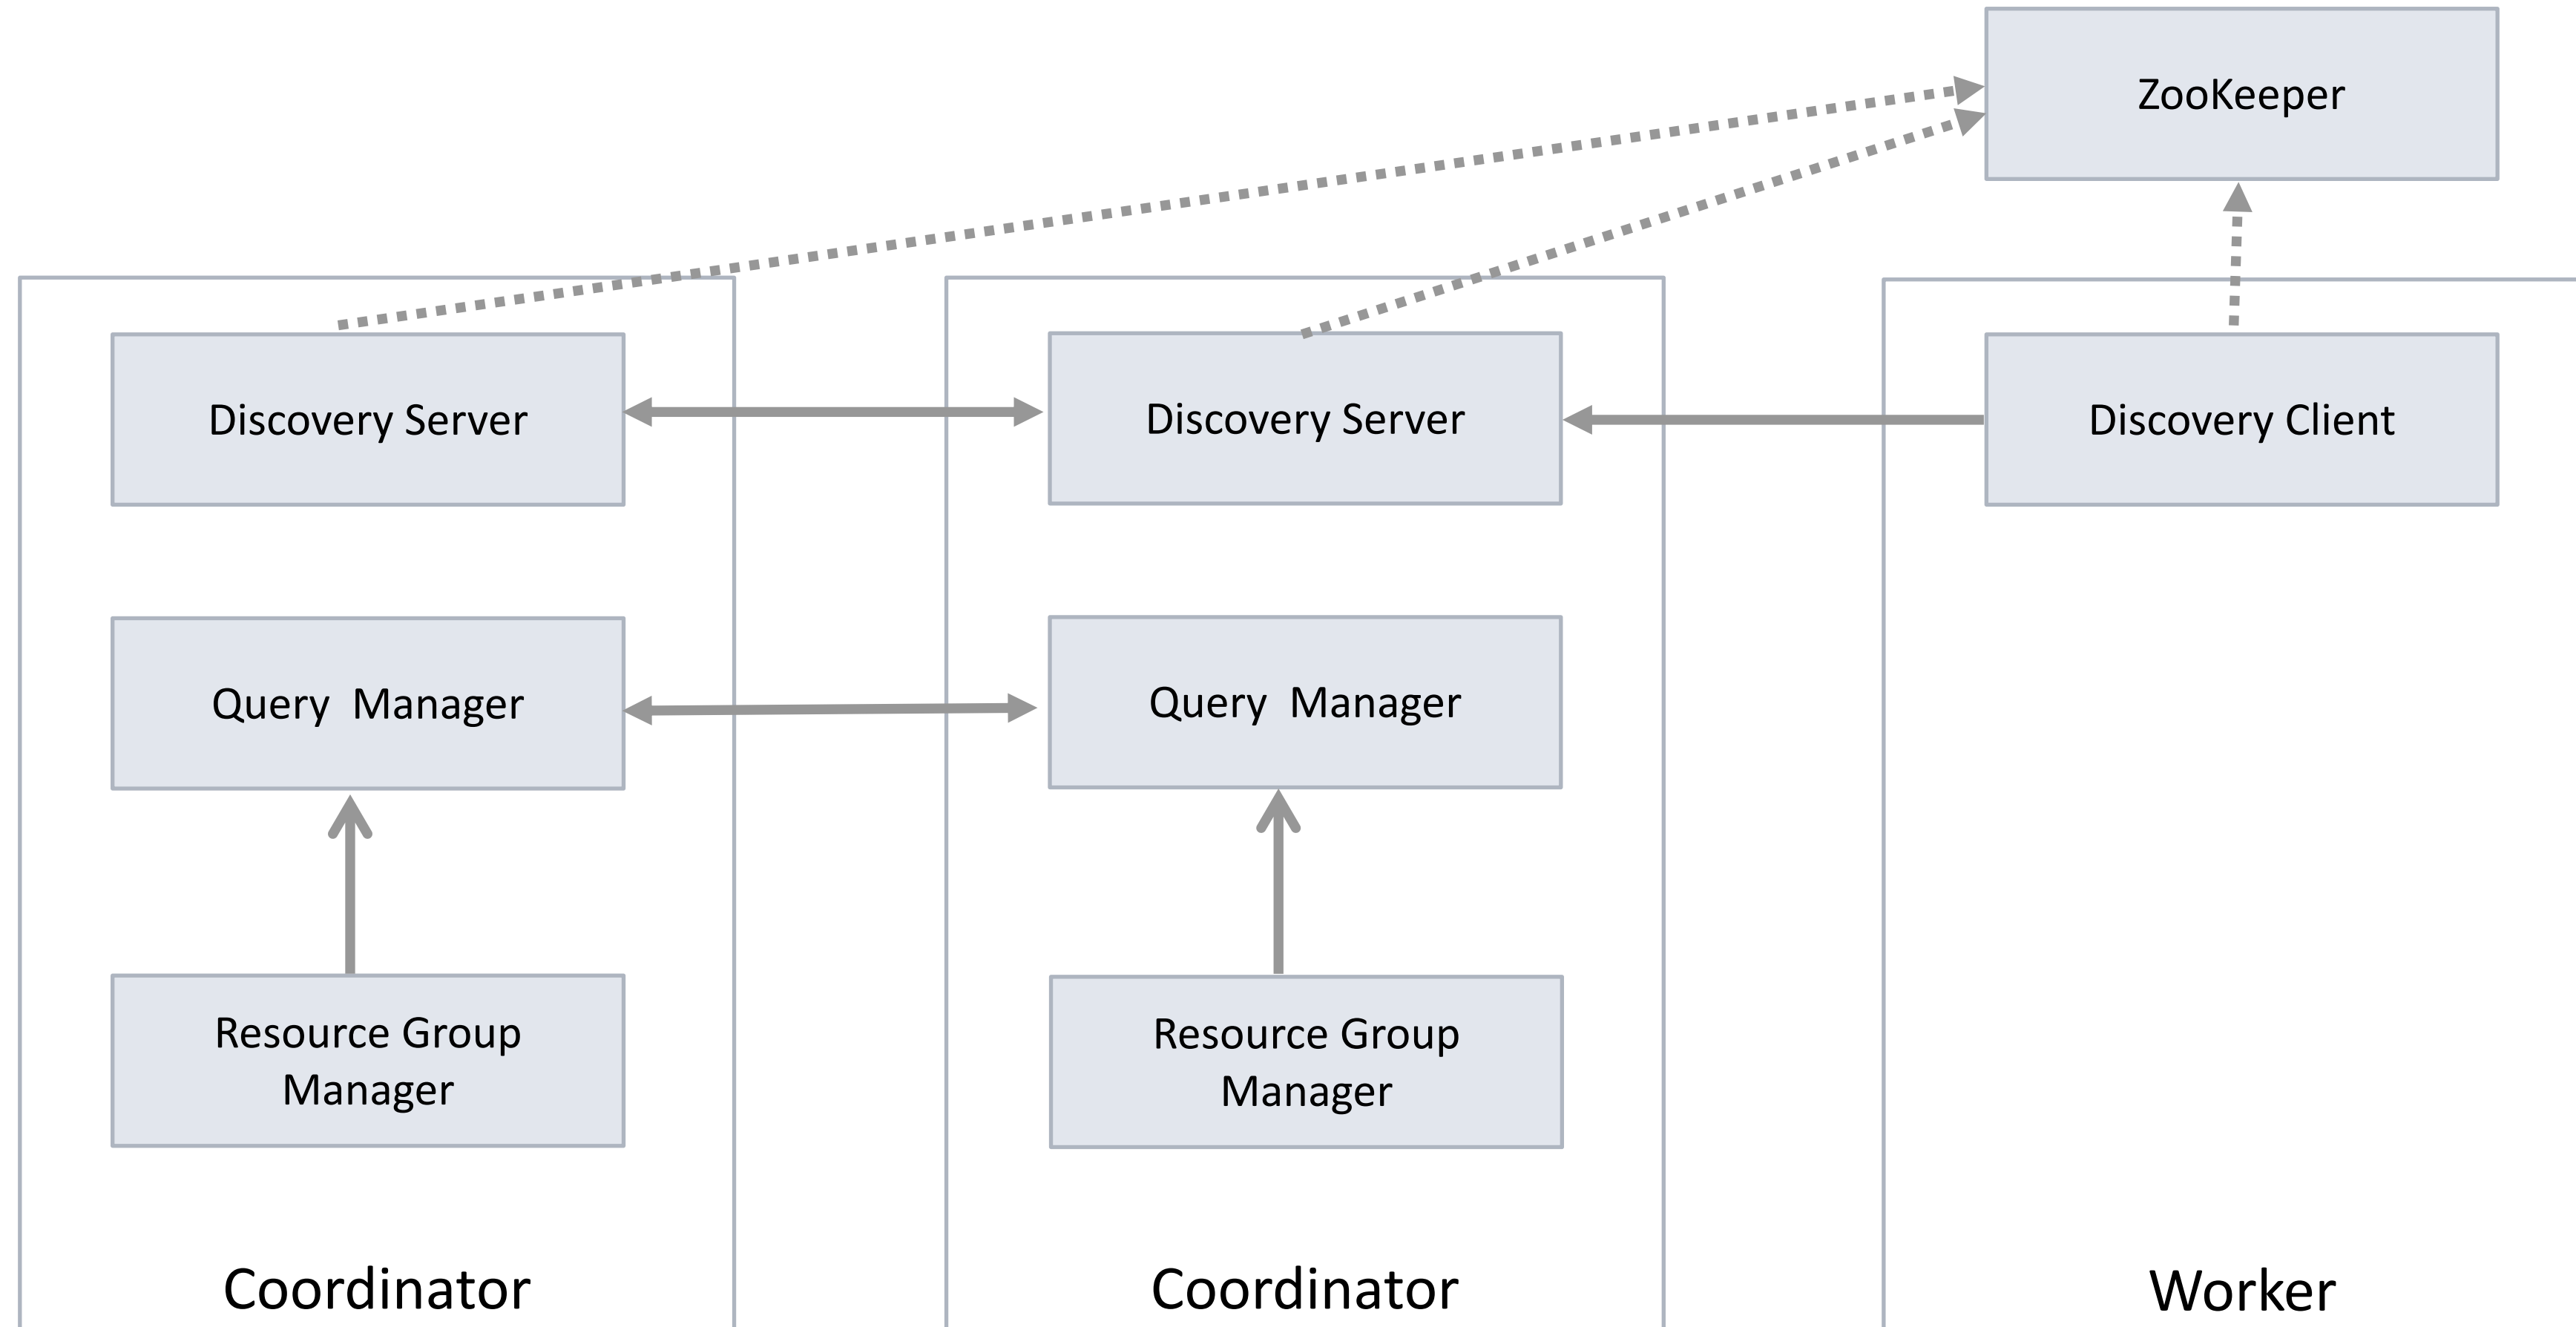
2

Presto Cluster Stability and Maintainability Improvement

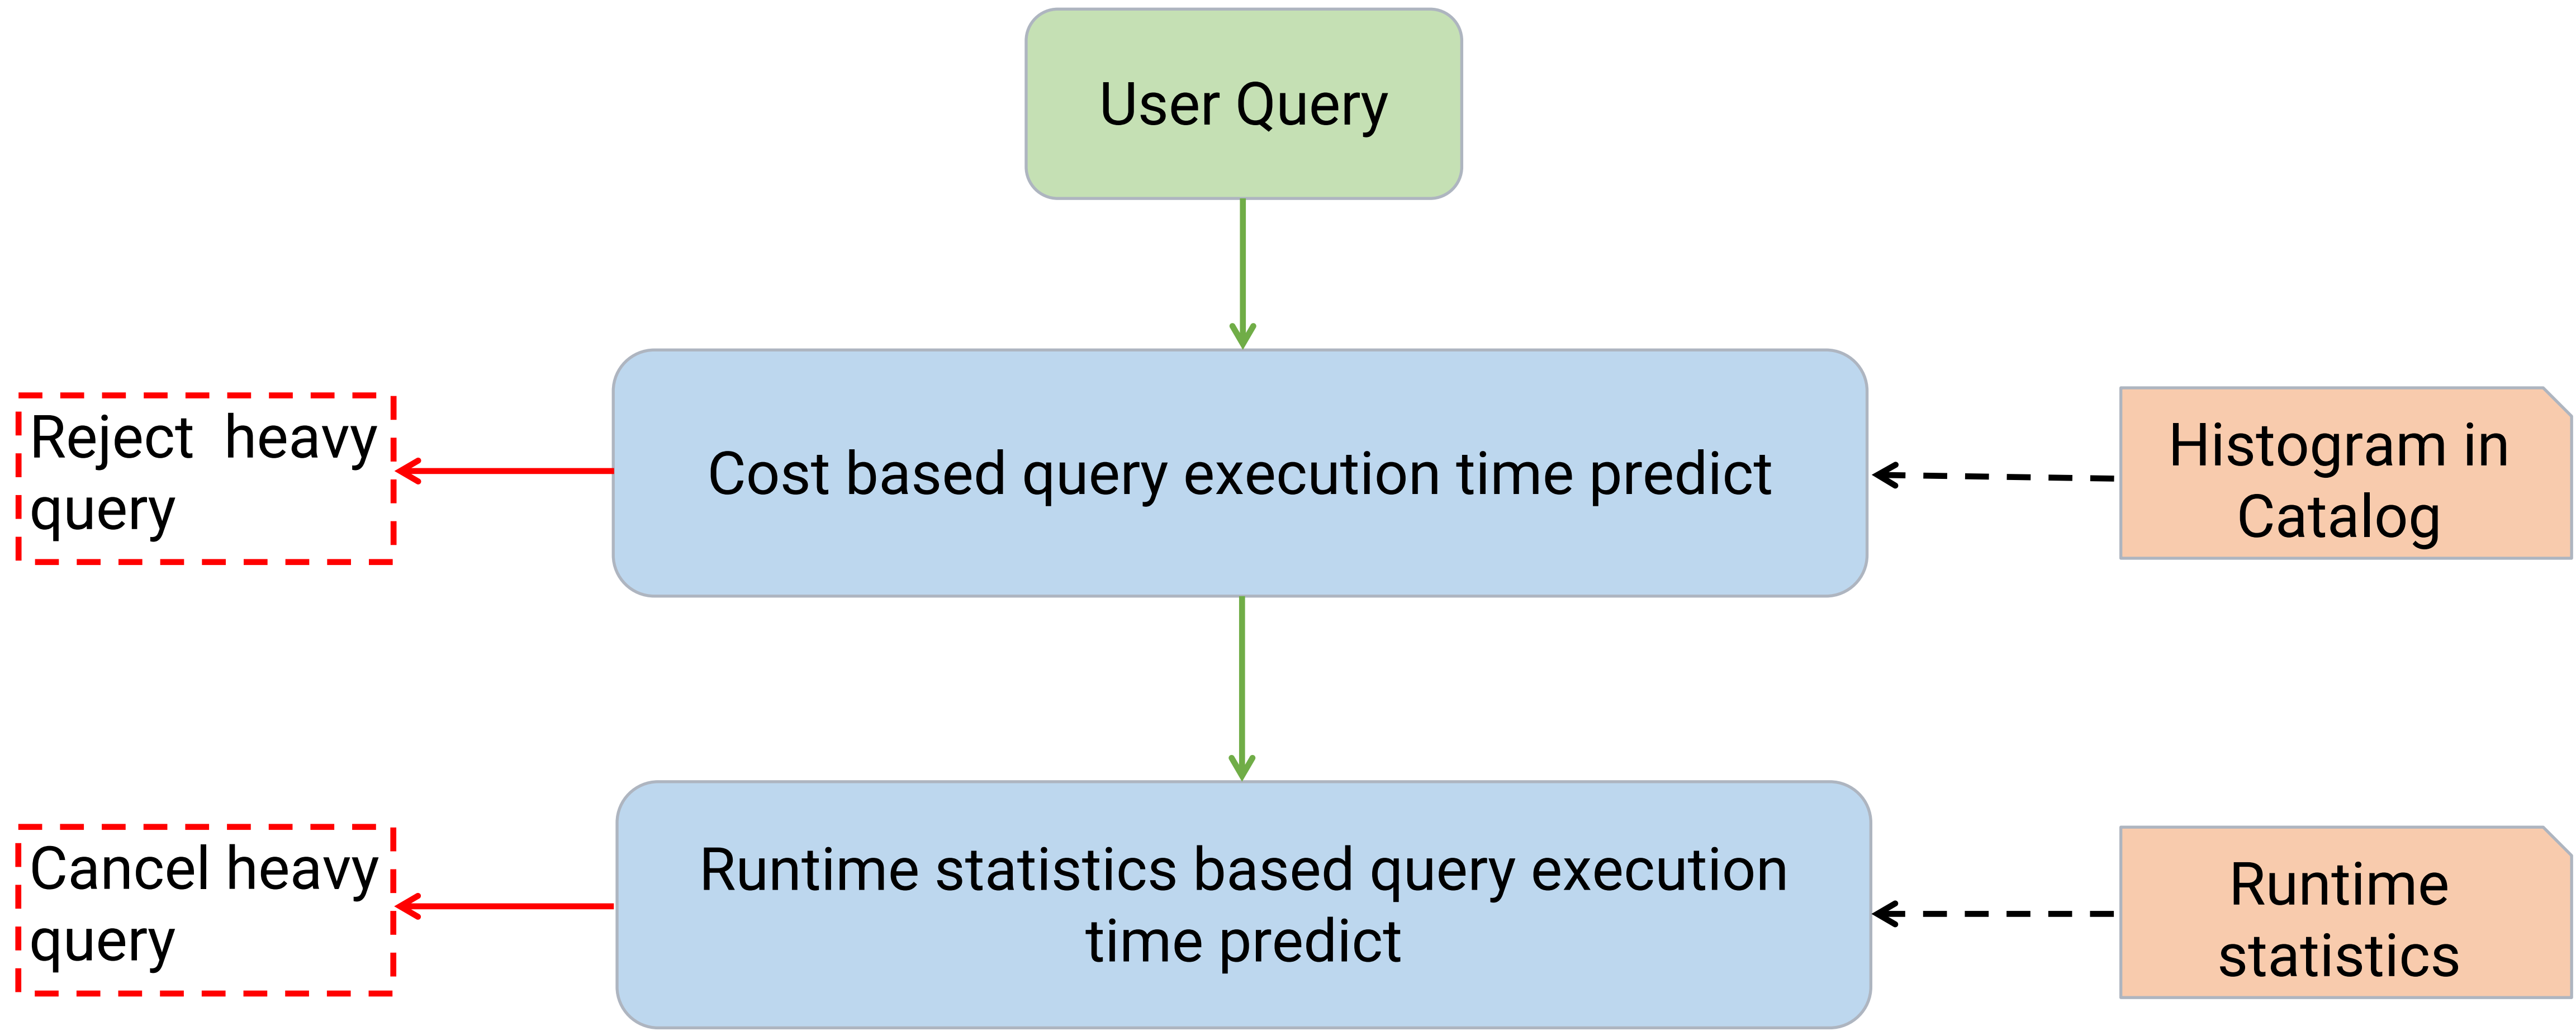
Gateway: routing and throttling



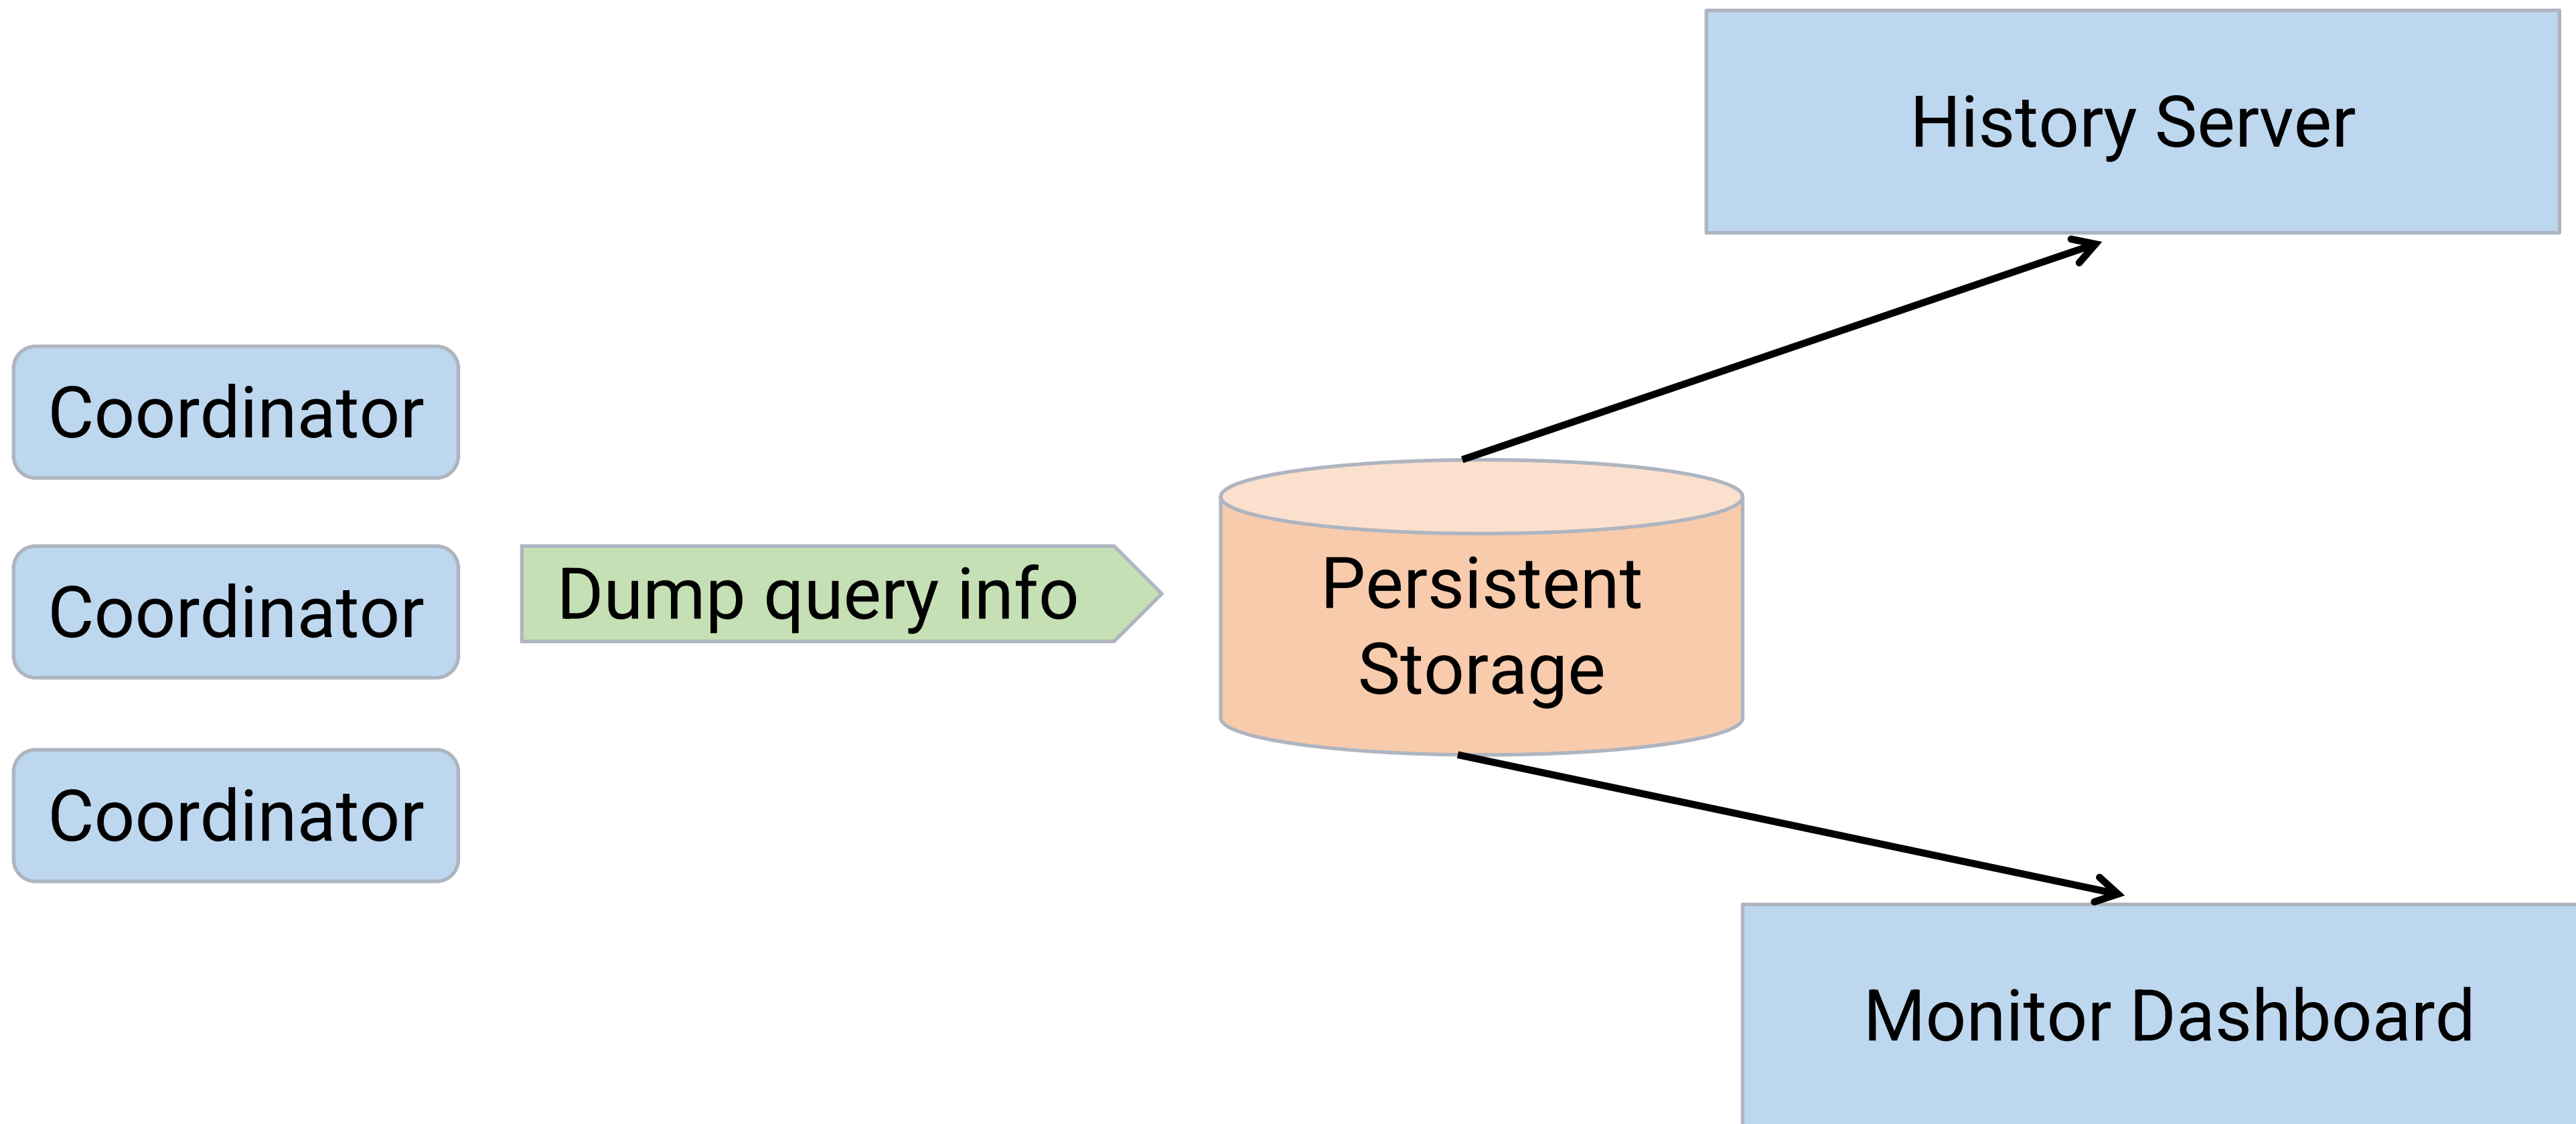
Multi-coordinator to ensure high availability



Query execution time predict and adaptive cancellation to avoid heavy query



History Server: persistent query information for history analysis

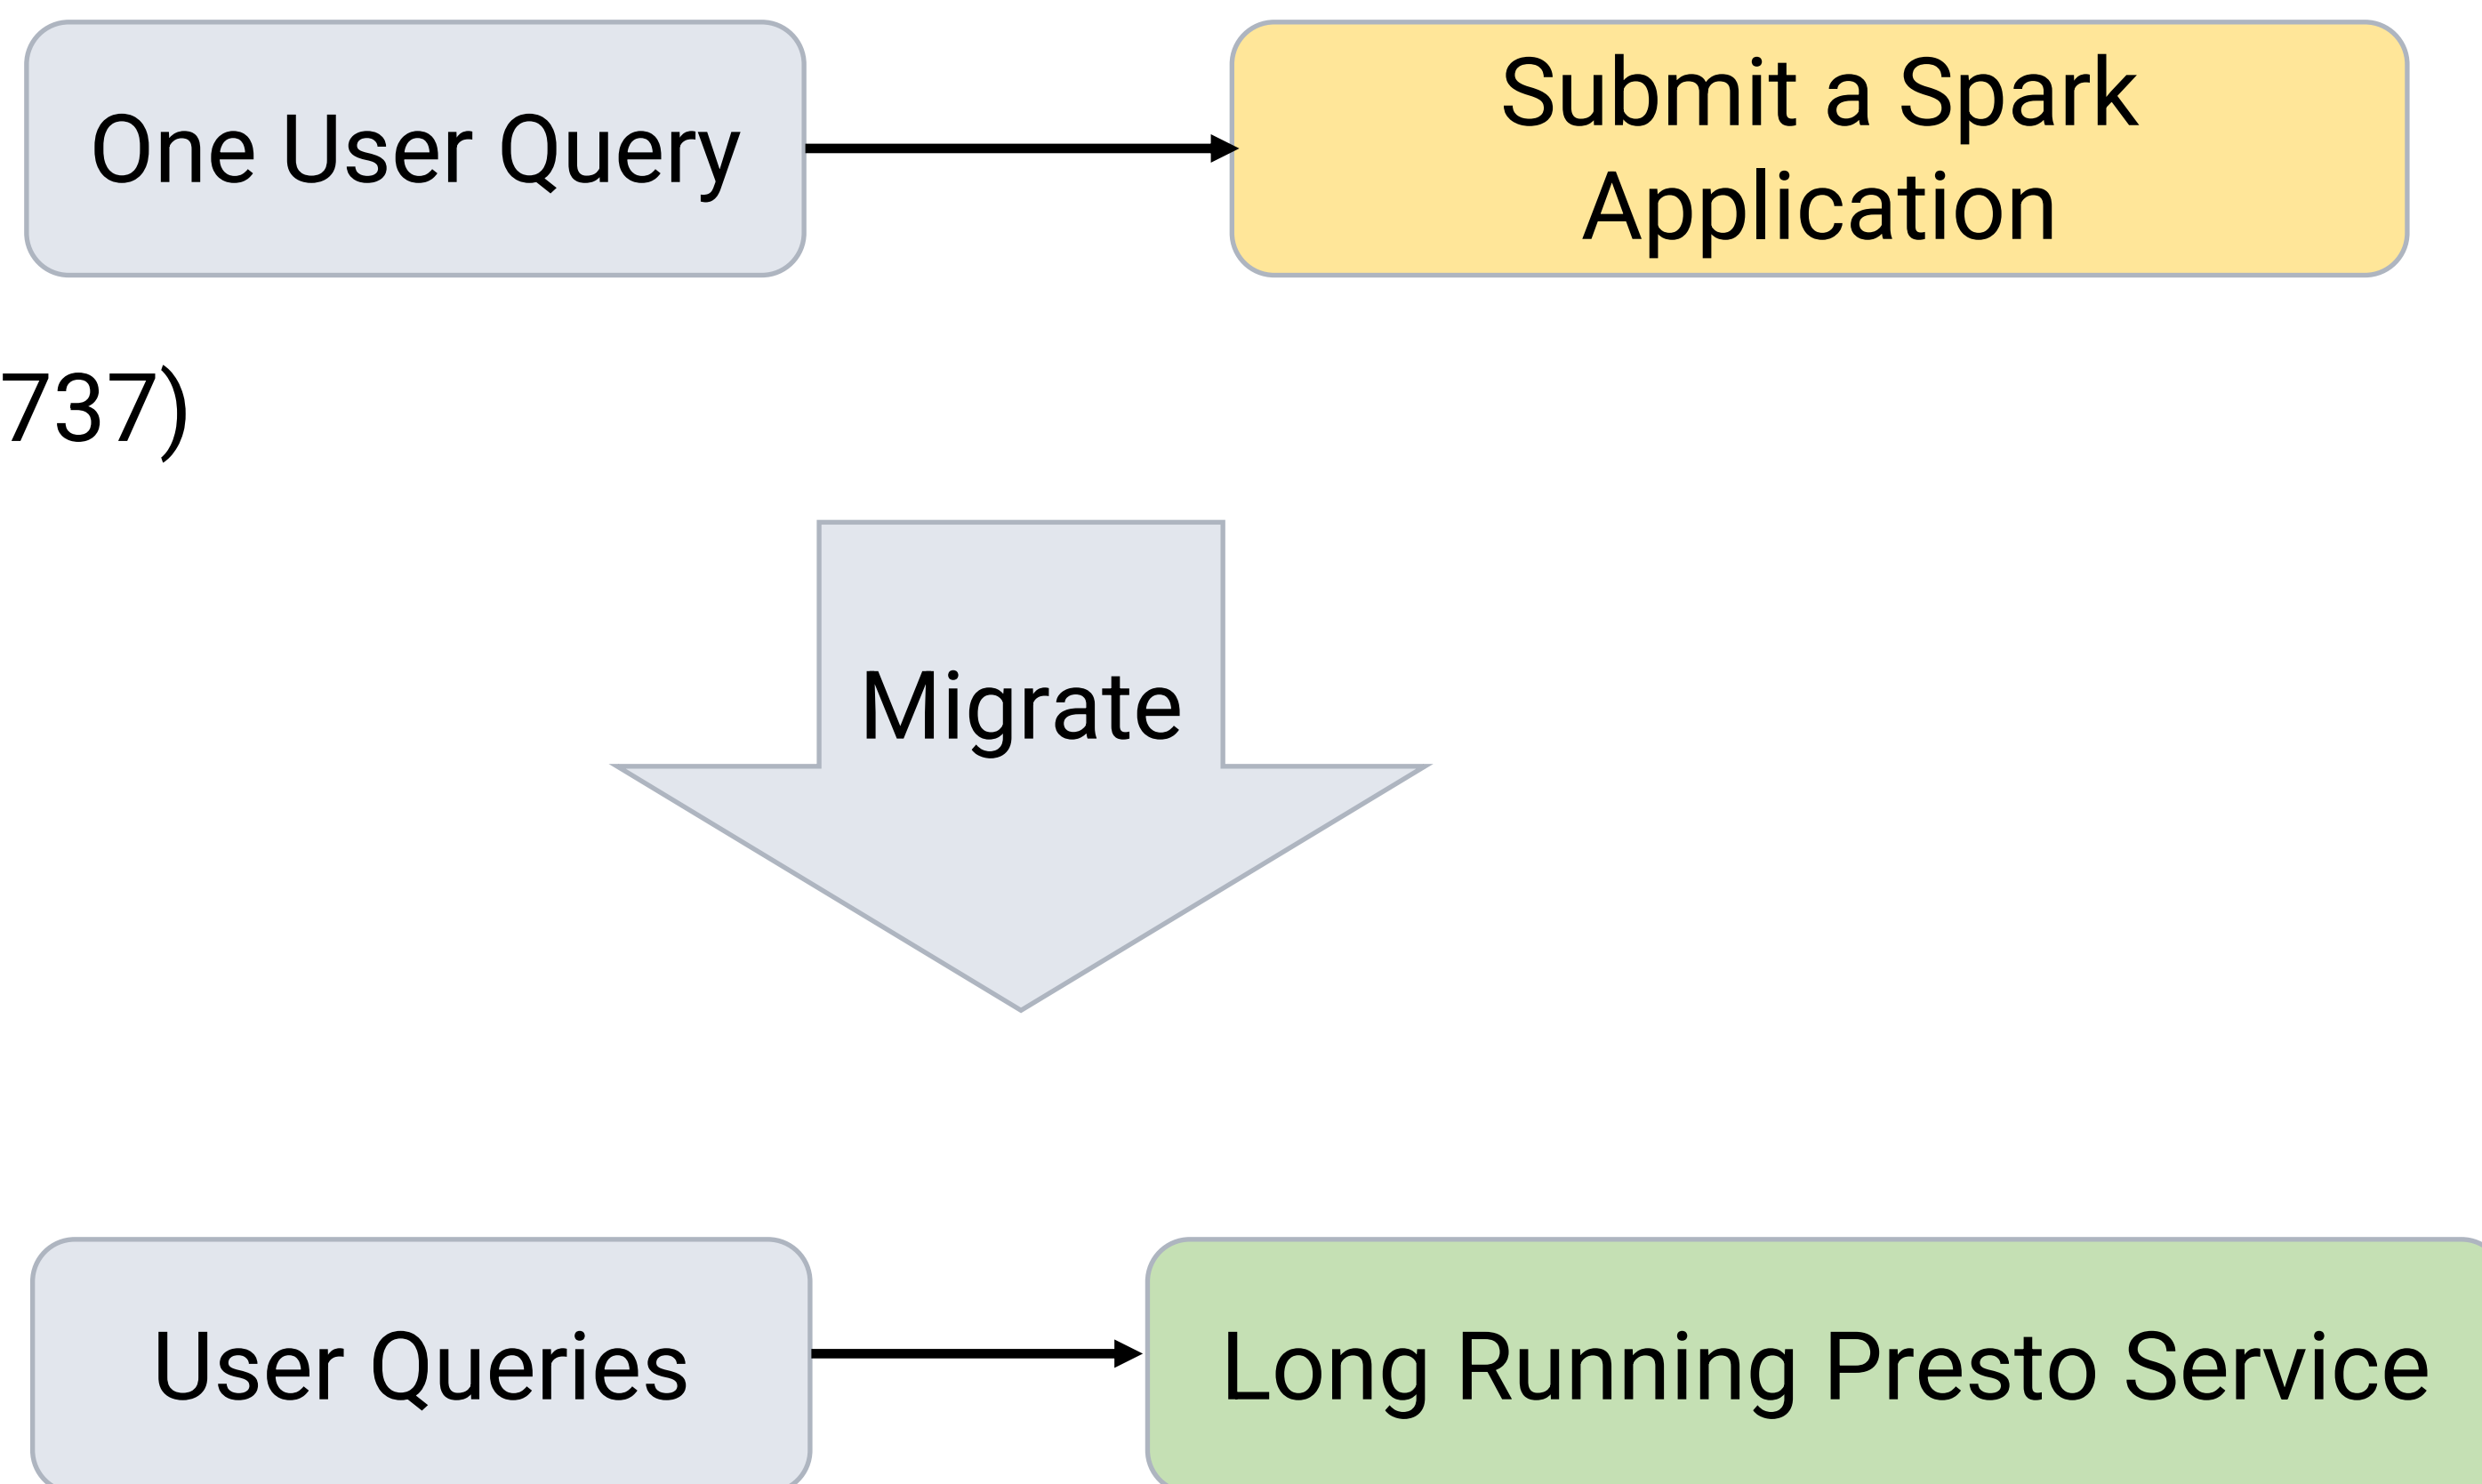


3

Practice and Optimization at Different Scenarios


Accelerate Ad-hoc queries by migrating from Hive/SparkSQL to Presto

- Hive UDF Wrapper
(<https://github.com/prestodb/presto/pull/16737>)
- Support range partition
- Hive semantics compatible
array index, divide by zero etc.
- Support new syntaxs:
 - Insert HDFS directly
 - Extend Inserting hive syntax
 - etc.



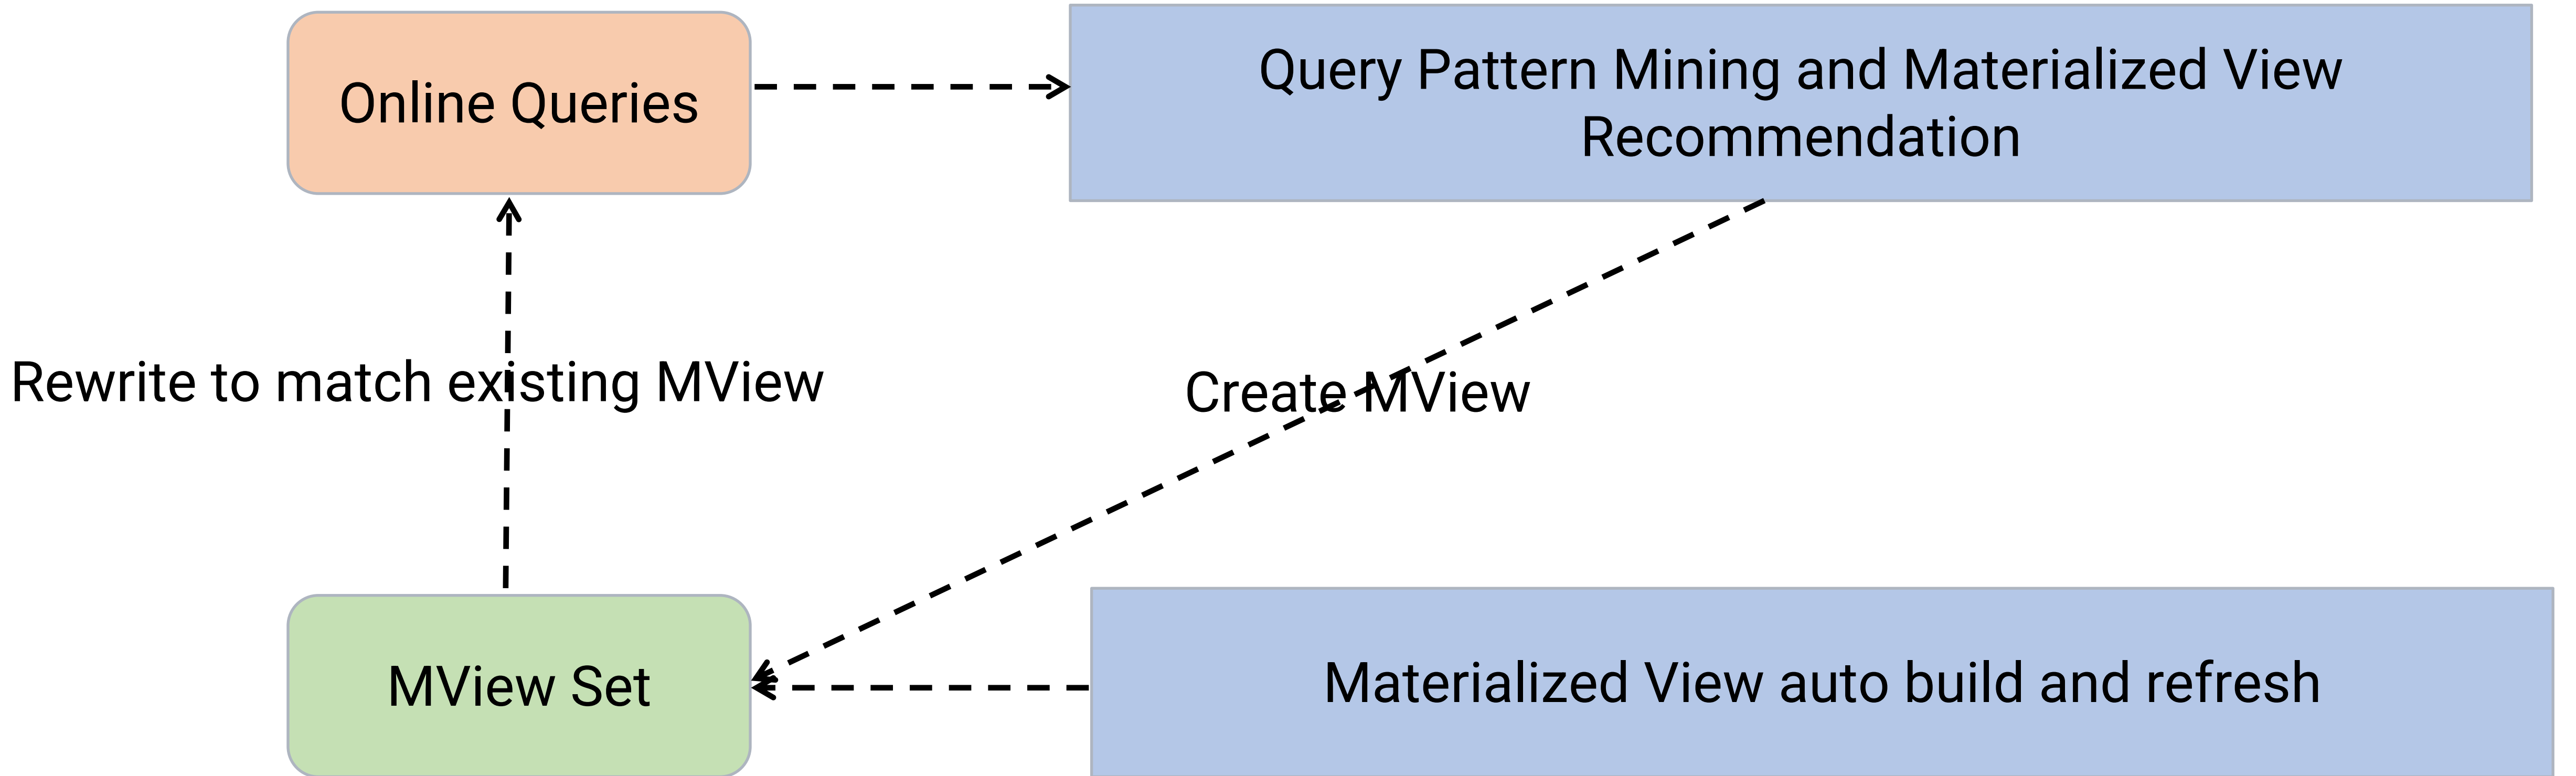
Accelerate BI analysis by materialized view

- Generate SQL by template, pattern is fixed
- High QPS, Low Latency
- Drill down / Roll up / Filter on different dimensions using same dataset



BI Platform Example

Accelerate BI analysis by materialized view



Accelerate BI analysis by materialized view

Materialized view:

```
1 CREATE MATERIALIZED VIEW mv1
2 AS
3 SELECT empid
4 FROM emps JOIN depts ON emps.deptno = depts.deptno;
```

Query:

```
1 SELECT empid
2 FROM depts JOIN (
3 SELECT empid, deptno
4 FROM emps
5 WHERE empid = 1) AS subq
6 ON depts.deptno = subq.deptno
```

Rewriting:

```
1 SELECT empid
2 FROM mv1
3 WHERE empid = 1
```

Join Rewrite

Materialized view:

```
1 CREATE MATERIALIZED VIEW mv3
2 AS
3 SELECT empid, deptno, COUNT(*) AS c, SUM(salary) AS s
4 FROM emps
5 WHERE deptno > 5
6 GROUP BY empid, deptno
```

Query:

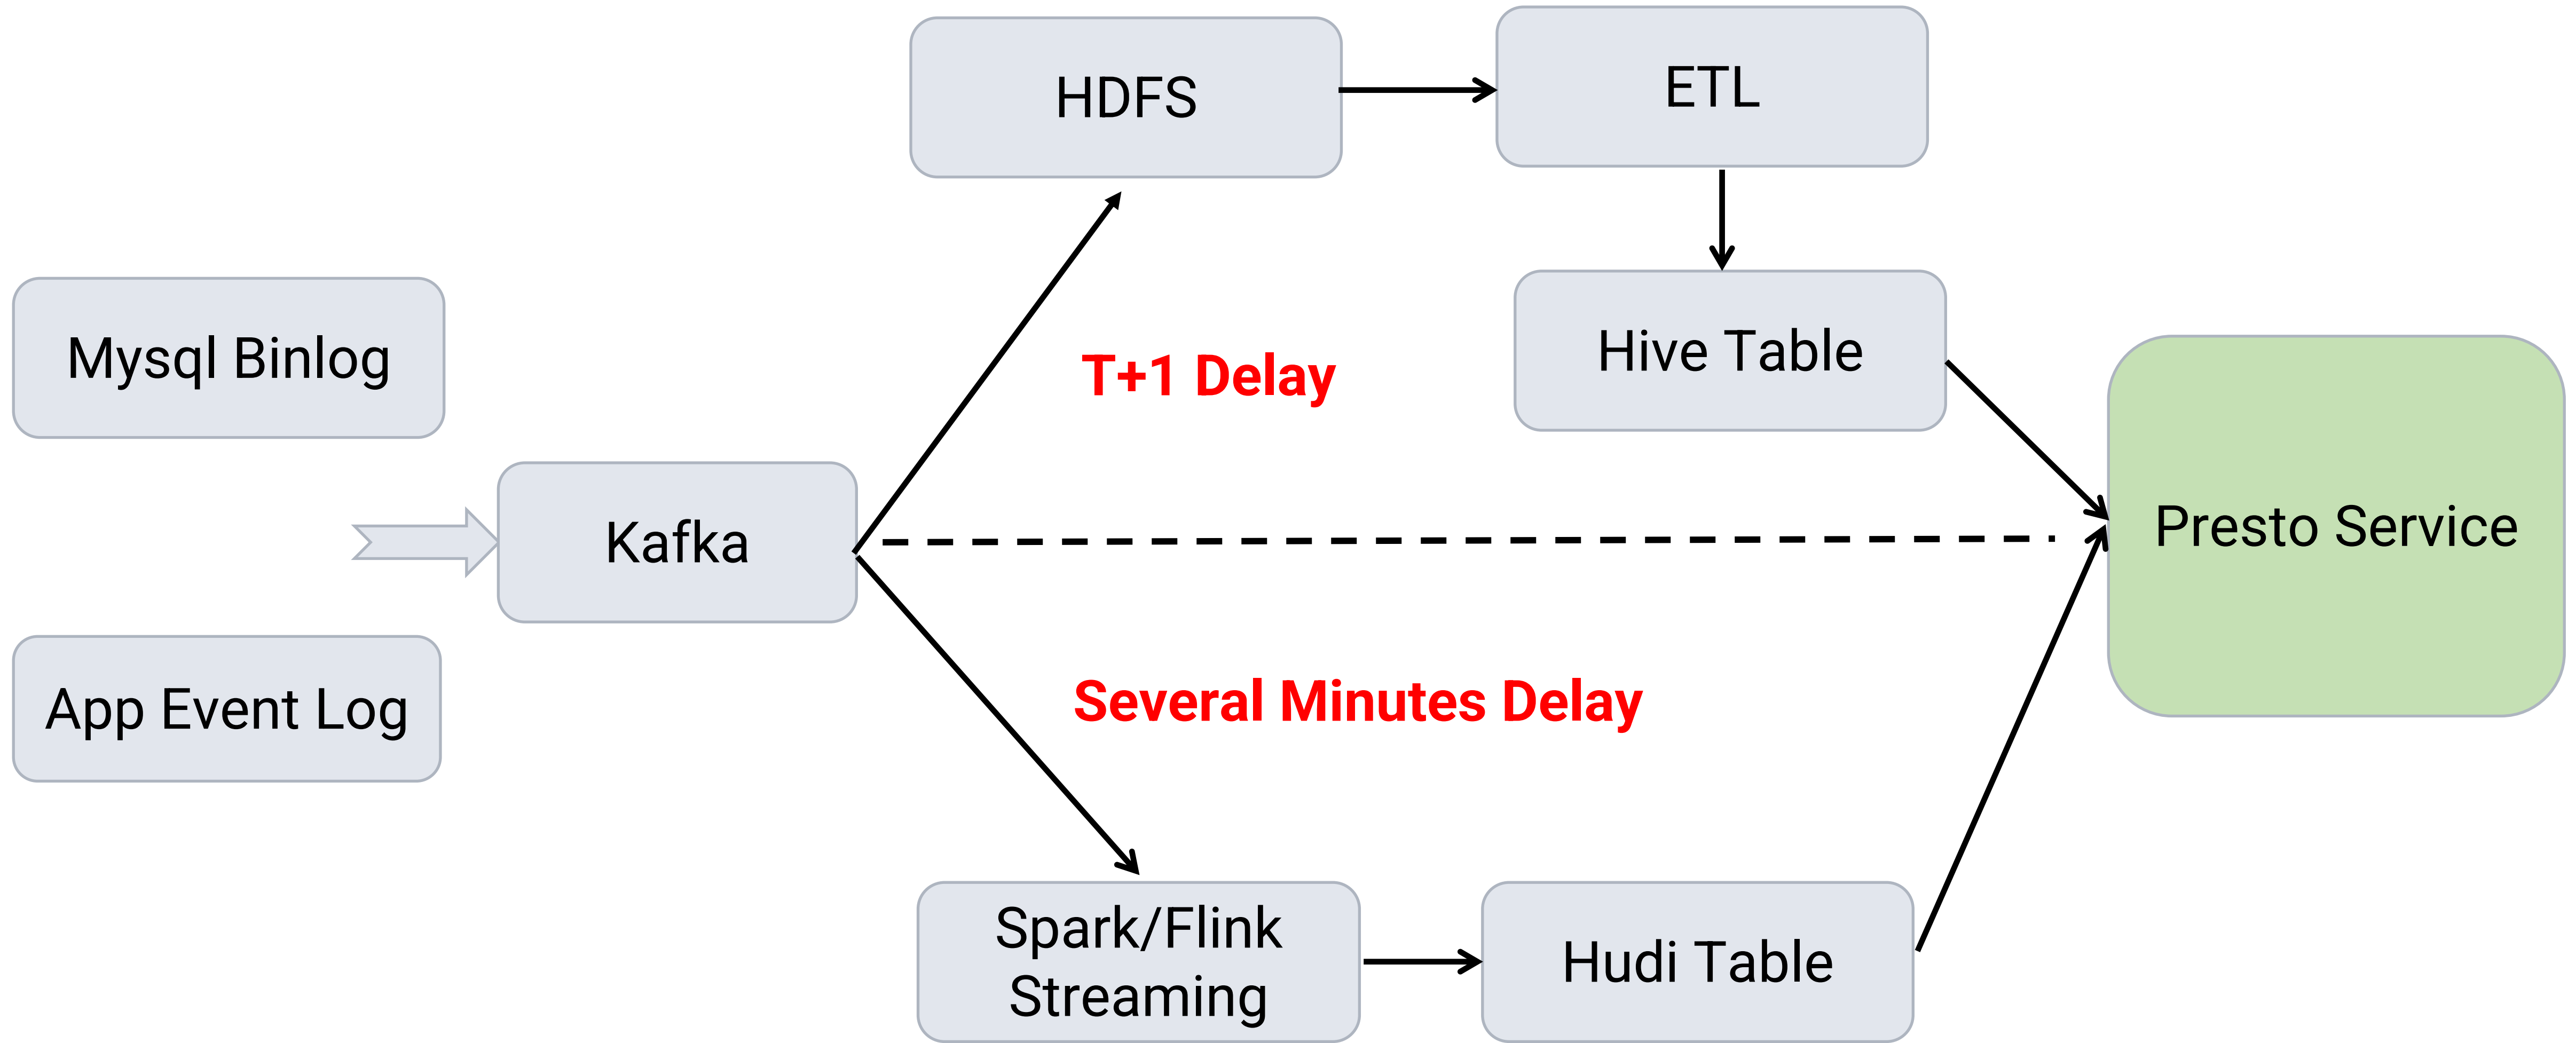
```
1 SELECT deptno, COUNT(*) AS c, SUM(salary) AS s
2 FROM emps
3 WHERE deptno > 10
4 GROUP BY deptno
```

Rewriting:

```
1 SELECT deptno, SUM(c) AS c, SUM(s) AS s
2 FROM mv3
3 WHERE deptno > 10
4 GROUP BY deptno
```

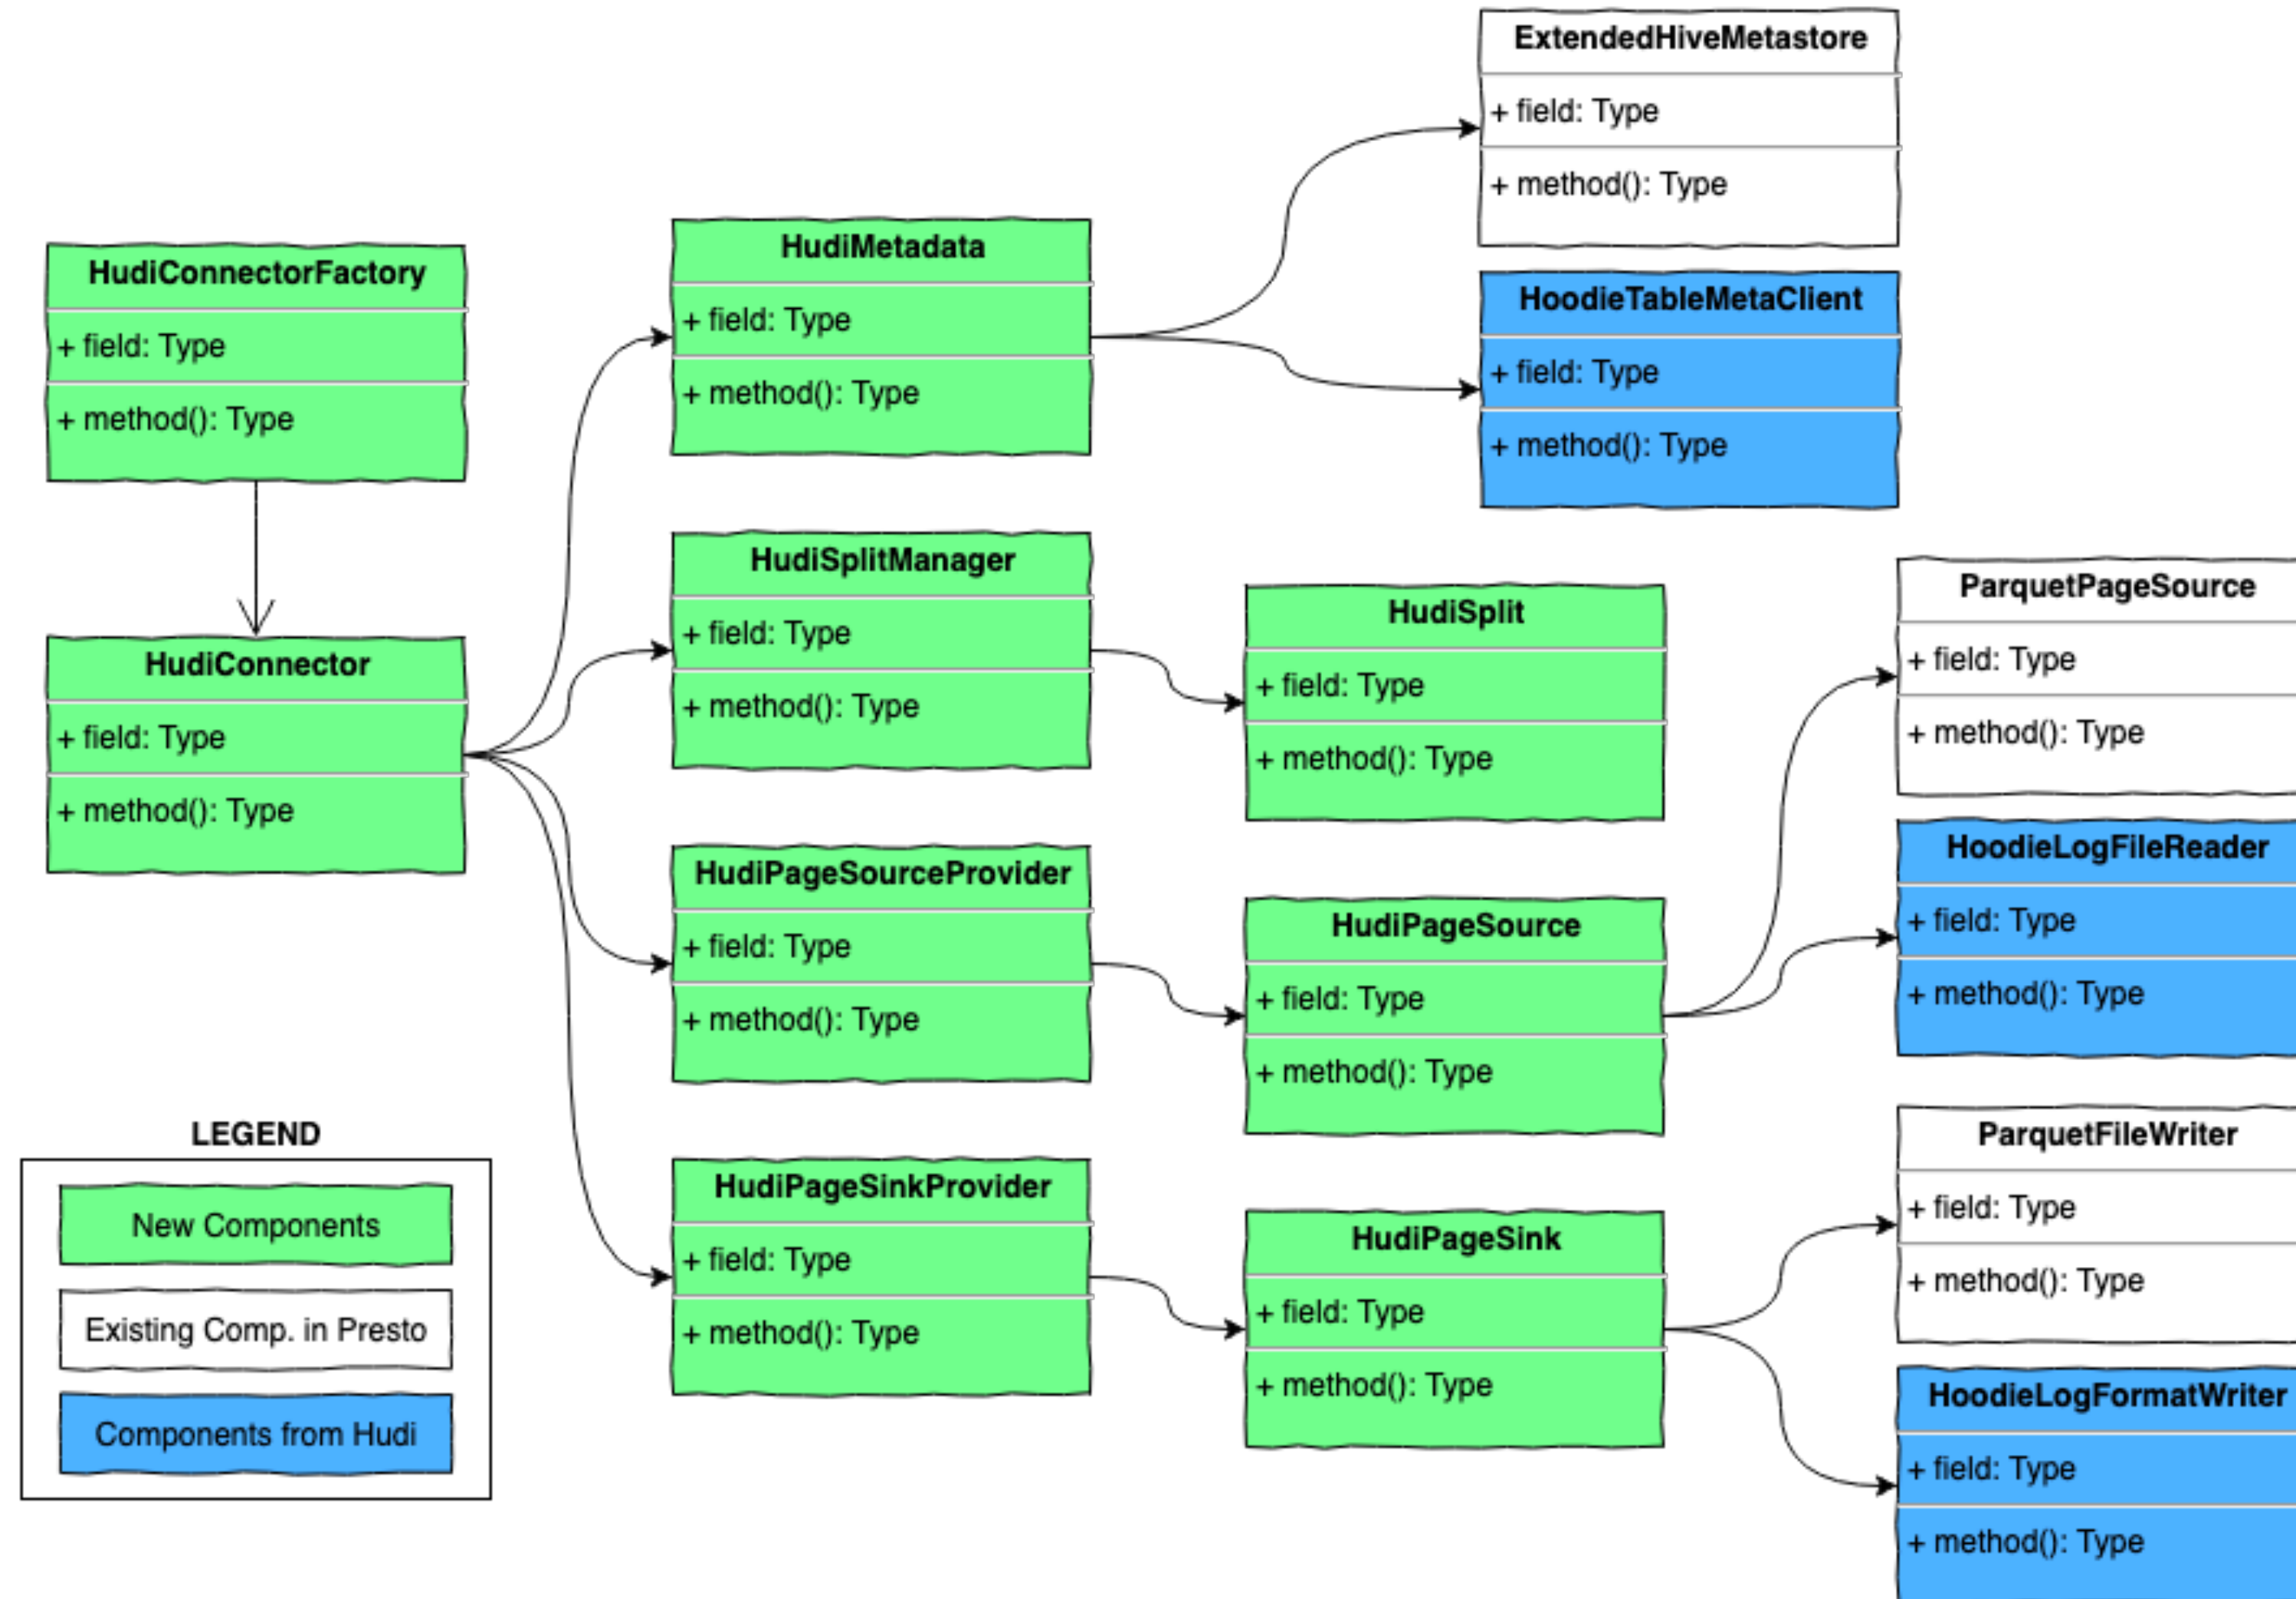
Aggregate Rewrite

Enable near real-time data analysis by support hudi connector



Enable near real-time data analysis by support hudi connector

Proposal in Community: <https://github.com/prestodb/presto/issues/17006>



4

Ongoing Exploration and Future Plan

- Serverless Presto on Public Cloud (<https://www.volcengine.com/product/las>)
- Native Engine

*Inspire
Creativity
Enrich
Life*

 ByteDance 字节跳动

Thanks & QA