

DATA+AI
SUMMIT 2022

Interactive Querying with Spark SQL

ORGANIZED BY  databricks



Ashish Singh

Tech Lead, Big Data Query Platform,
Pinterest

Querying

Extracting information from data using SQL.

Querying

What about querying metadata of data?

Querying

- **Scheduled Query:** Query NOT executed by a user.

Querying

- **Scheduled Query:** Query NOT executed by a user.
- **Interactive Query:** Query executed by a user.

Scheduled vs Interactive Querying

	Scheduled	Interactive
Predictability	High	None
Results	Based on SLAs	Immediate
Syntax/Analysis errors	Low	High
ROI on Tuning	High	Low
Resource Efficiency Importance	High	Low

Data Scale at Pinterest

- 700+ PB in AWS S3
- 150K+ Data Compute Jobs per day (not including metadata only operations)
- ~20K Hadoop Nodes
- 1000+ Presto Workers
- 300K+ Hive Tables
- Everything in Cloud (AWS)

Query Engines at Pinterest

- Spark SQL
- Presto
- Flink SQL

Learn more about our Hive to Spark SQL migration

Advanced Migrations: From Hive to Spark SQL

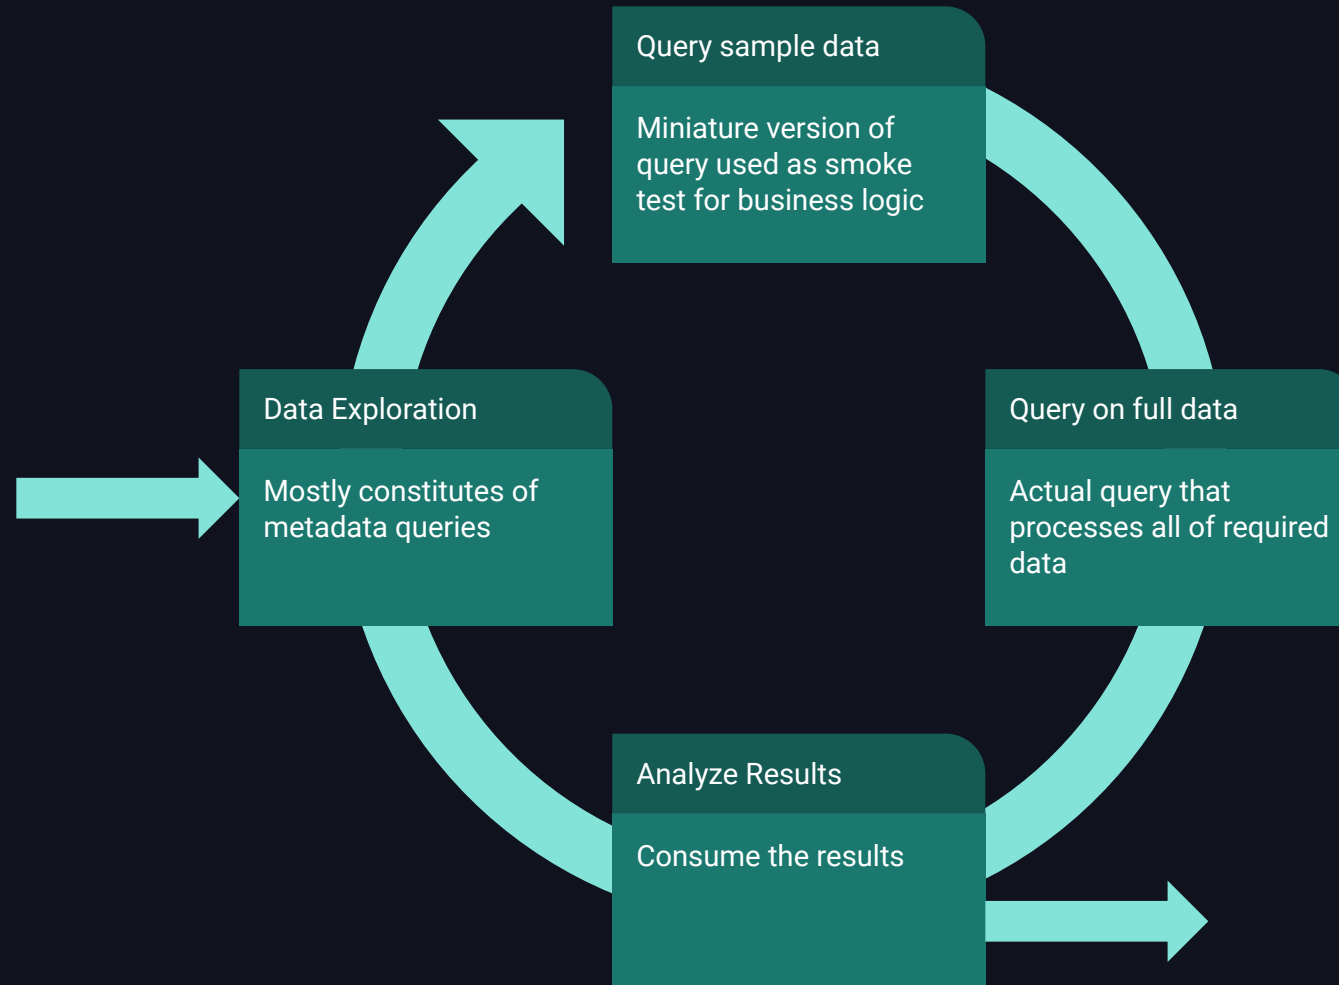
Wednesday, June 29 @4:45 PM

MOSCONE SOUTH | UPPER MEZZANINE | 211

By Zaheen Aziz

Interactive Querying

Usage pattern

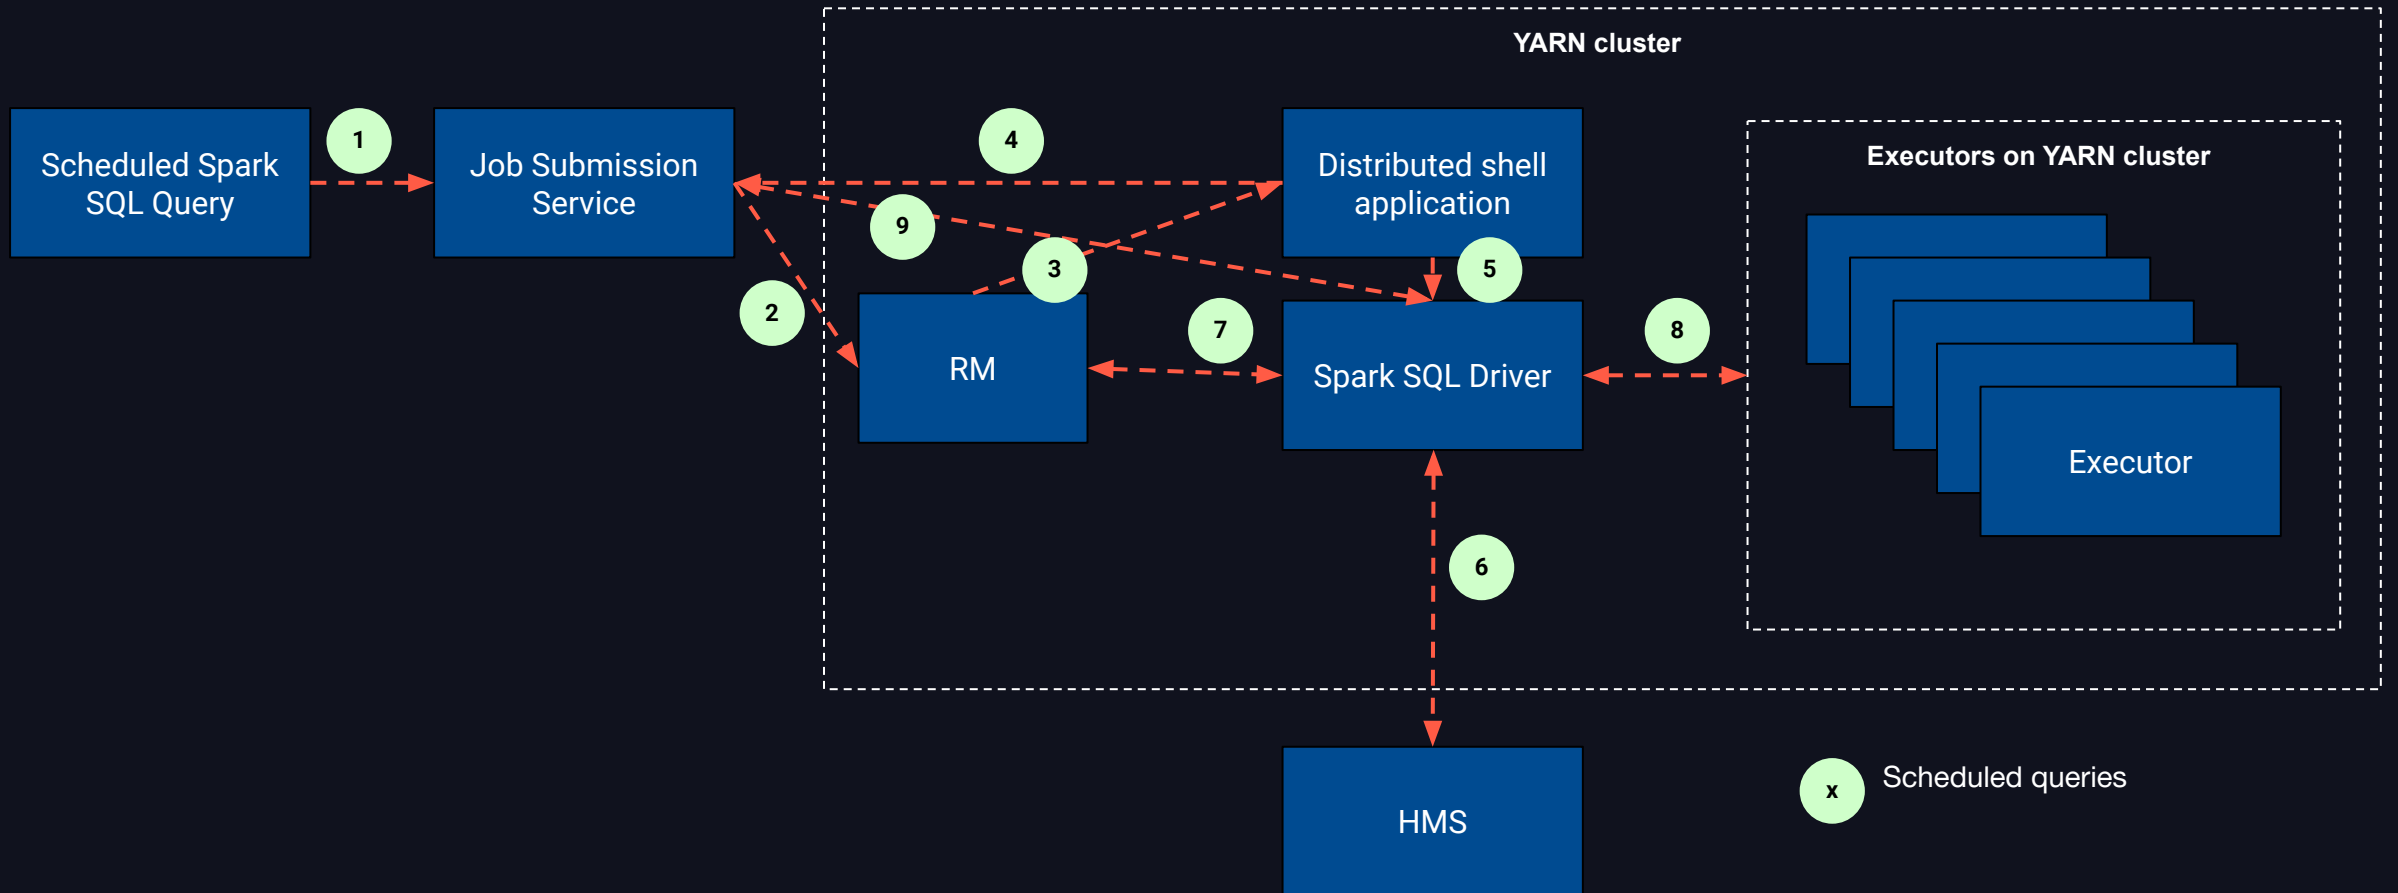


Interactive Querying

Requirements

- Seamless query submission
- Fast metadata queries
- Fail fast
- Easy access to results
- Error handling and tuning suggestions
- Auto error handling and retries

Scheduled Querying With Spark SQL



Interactive Querying

Requirements Check

- Seamless query submission ❌
- Fast metadata queries ❌
- Fail fast ❌
- Easy access to results ❌
- Error handling and tuning suggestions ❌
- Auto error handling and retries ❌

Architectural Choices

Distributed Shell Application

- Spark-sql CLI as dist shell application on YARN
- Waiting for container allocations
- Retrieving results is tricky
- Tracking statement-level progress is hard

Apache Spark Thrift Server

- Similar to HiveServer2
- JDBC/ ODBC protocols
- No isolation between queries

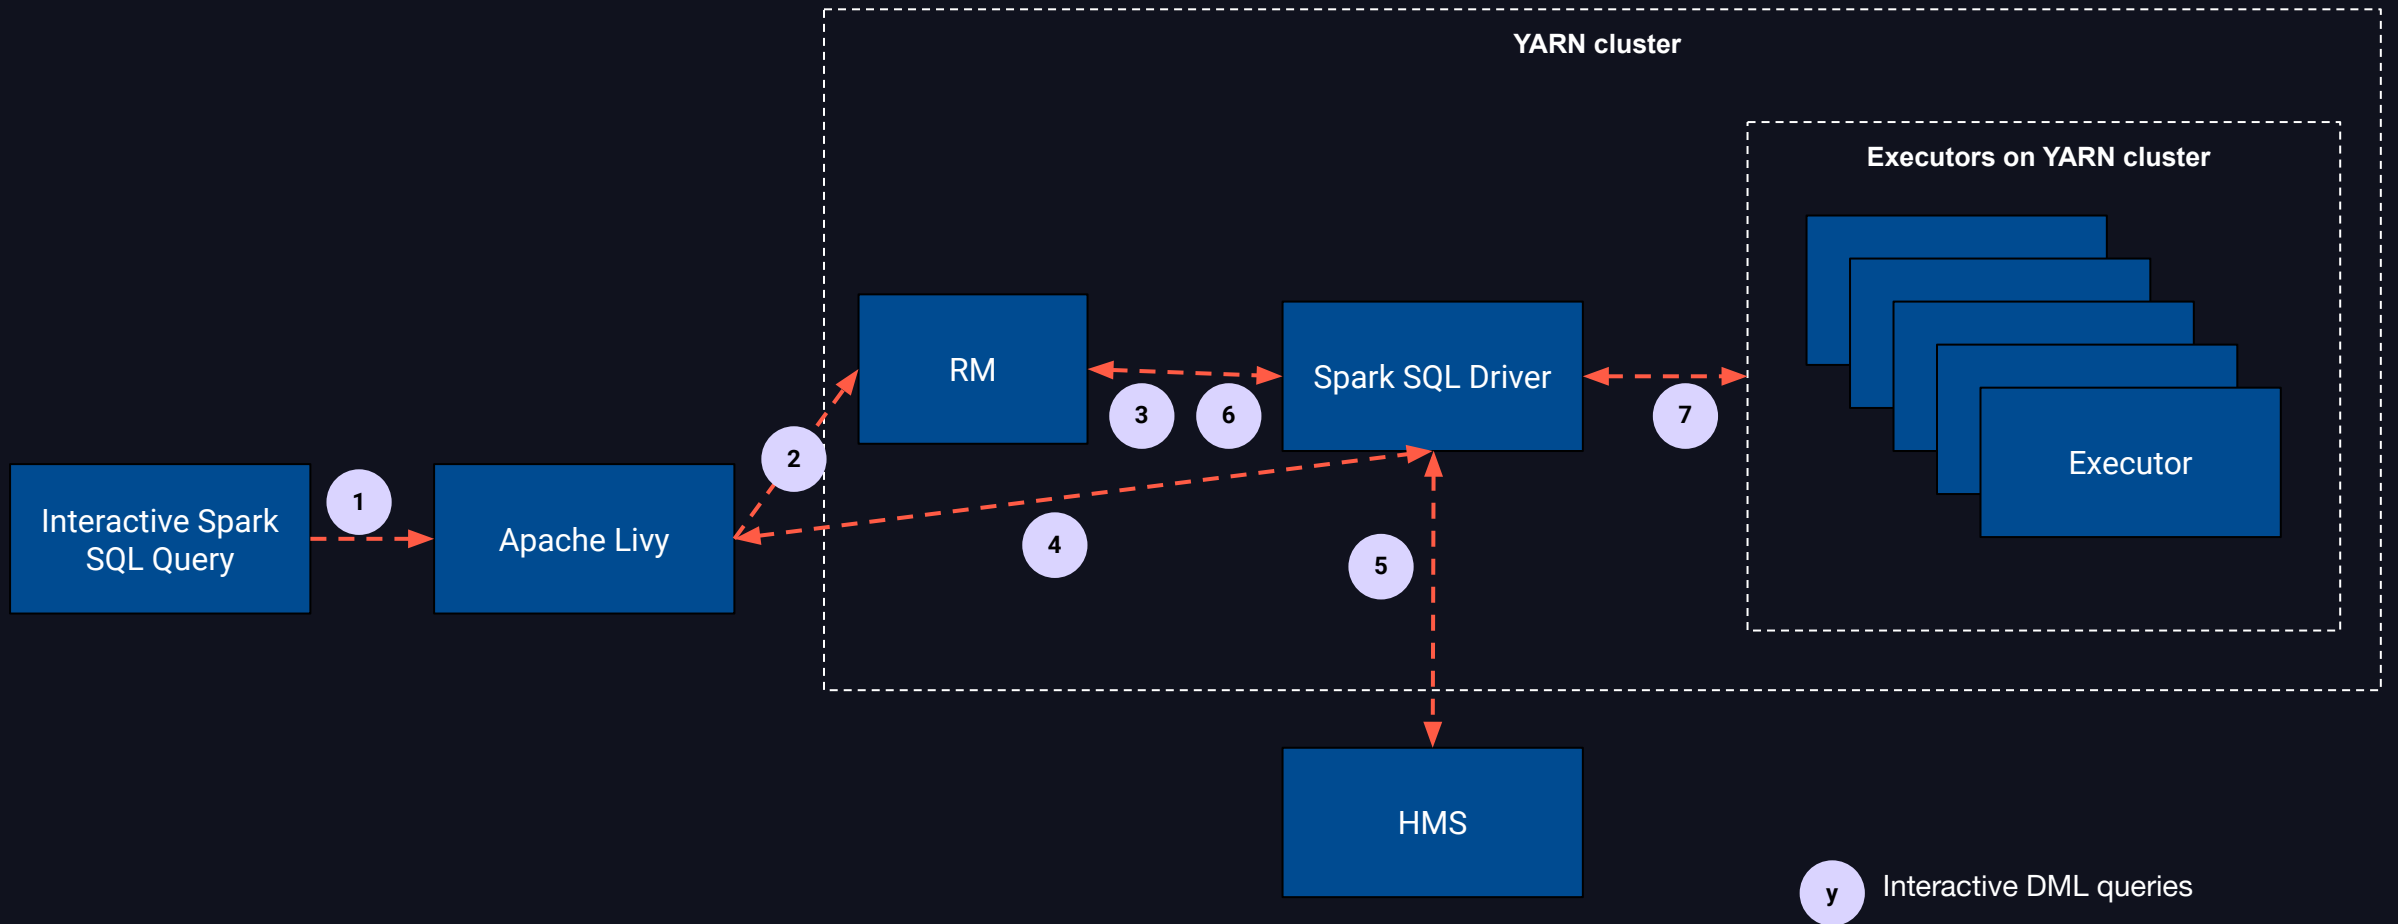
Apache Livy – Batch Sessions

- REST interface to Spark cluster
- Similar to spark-submit
- Multi-tenancy, high availability, and failure isolation

Apache Livy – Interactive Sessions

- Enables reusing sessions
- Multi-tenancy, high availability, and failure isolation

Querying With Spark SQL



Interactive Querying

Requirements Check

- Seamless query submission ❌
- Fast metadata queries ❌
- Fail fast ❌
- Easy access to results ❌
- Error handling and tuning suggestions ❌
- Auto error handling and retries ❌

Seamless Query Submission

From existing query clients like querybook, jupyter, etc

- Generic DB-API 2.0 compliant Python client on top of Livy's REST API

```
from bigpy.hive.livy.livy_client import Connection
from bigpy.hive.livy.livy_client import Cursor

cursor = Connection(livy_connection_urls='host1:port1').cursor()
cursor.execute('SELECT * FROM my_awesome_data LIMIT 10', async=True)

status = cursor.poll().operationState
while status == Cursor._STATE_RUNNING):
    logs = cursor.fetch_logs()
    for message in logs:
        print message
    status = cursor.poll().operationState

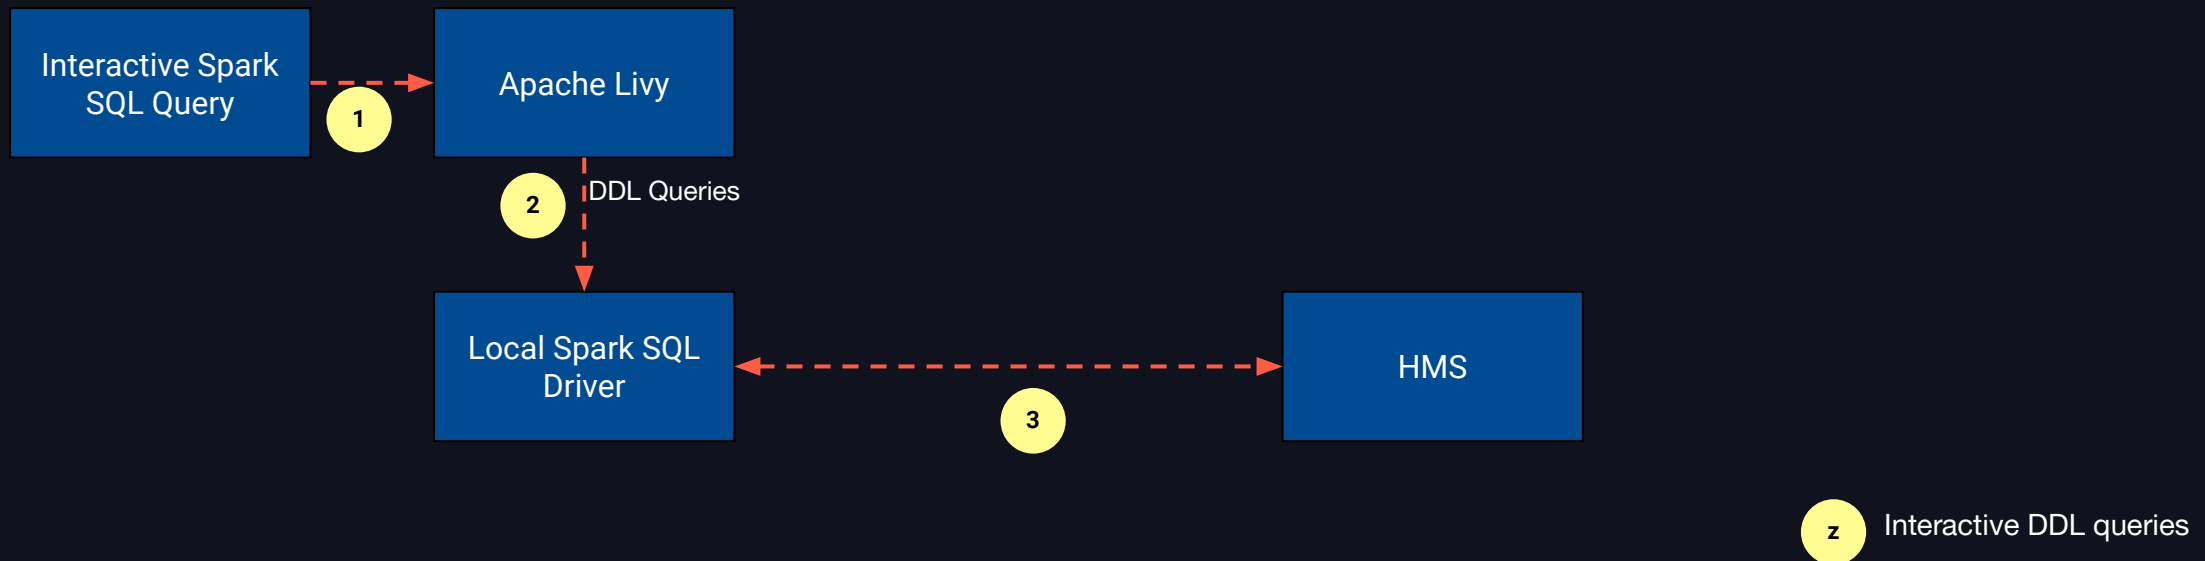
print cursor.fetchall()
```

Fast Metadata Queries

The problem..

- Waiting for container allocations
- Waiting for Spark Context to start up

Fast Metadata Queries

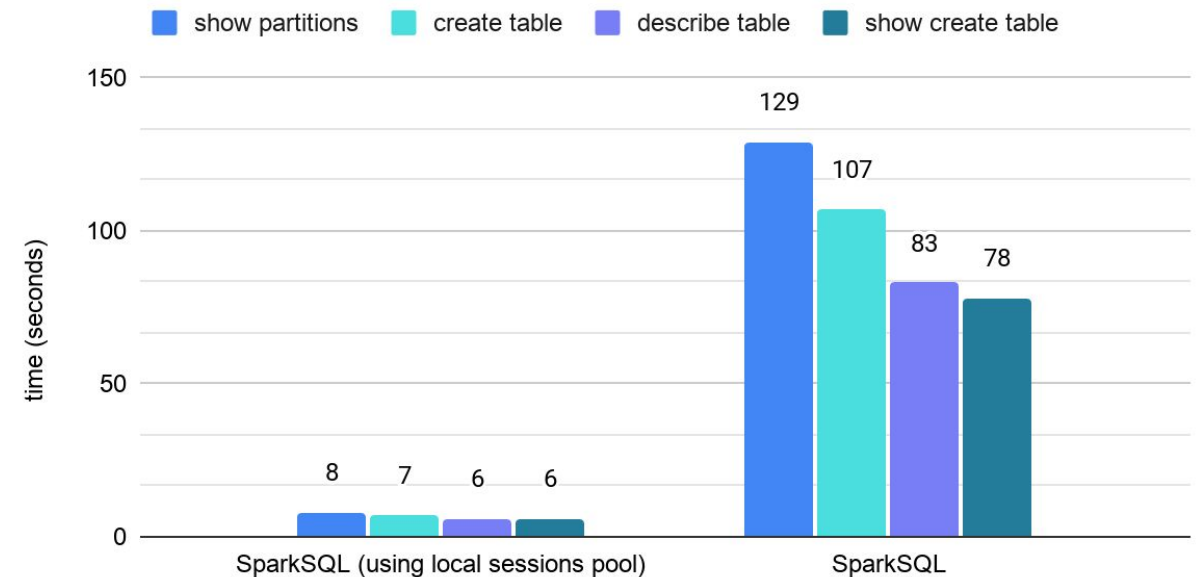


Fast Metadata Queries

- Pool is initialized on server startup
- Self-reliant, some of the features include:
 - Automatic garbage collection
 - Health monitoring
 - Asynchronous loading

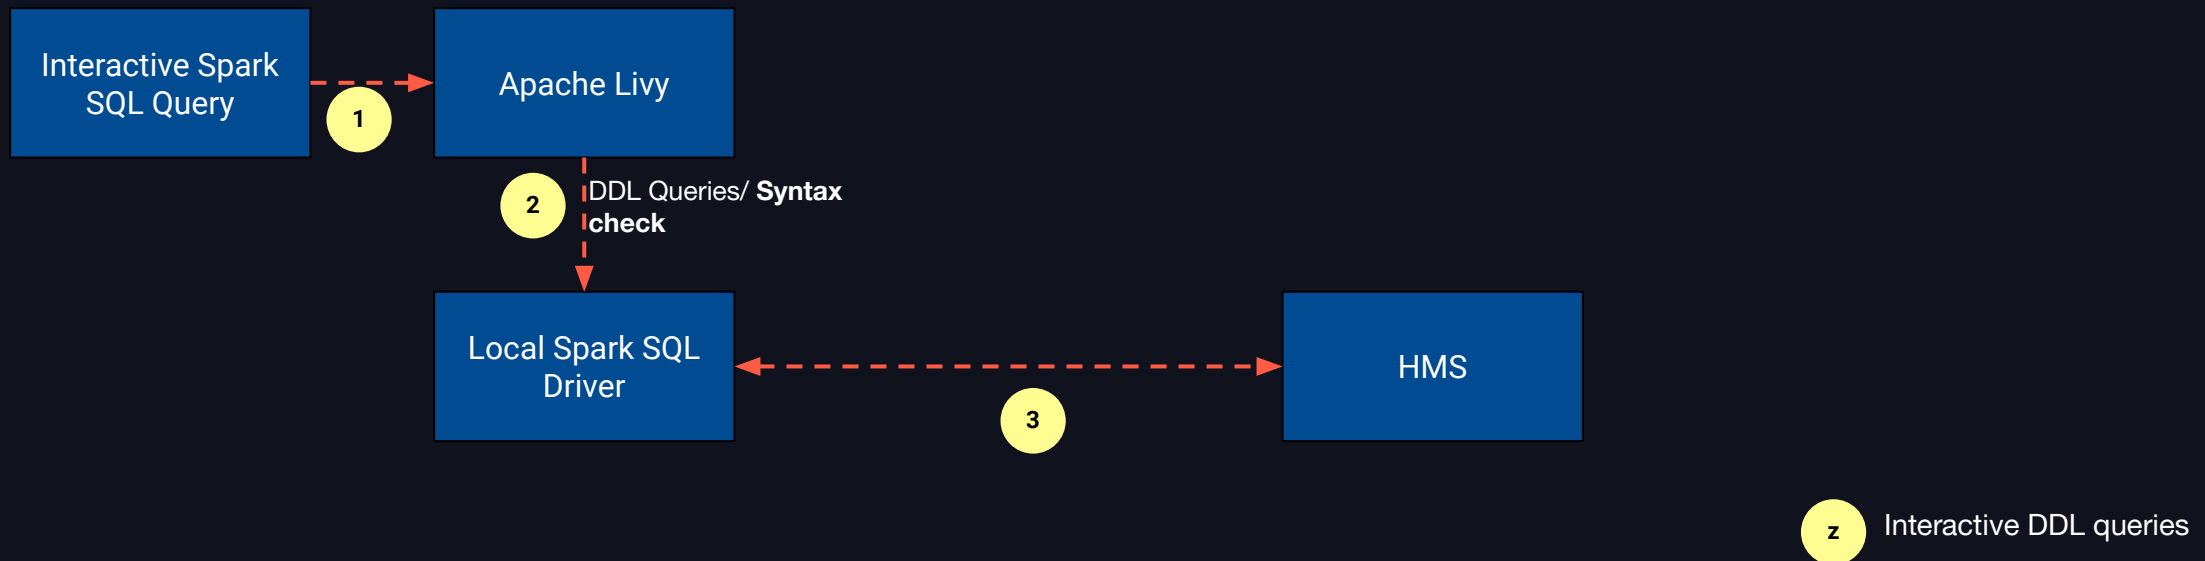
DDL query wall clock time comparison

Each datapoint is an average runtime of 30 past queries ran by the users



With this design, we reduced query latency from 70 seconds to an average of 10 seconds (~6.3x improvement)

Fail Fast



Easy Access to Results

- Livy limits the number of rows a query can return
- Needed to avoid Livy from choking on memory
- For large results, we redirect results to AWS S3
- Clients read results from S3 on query completion

Error Handling Recommendations

- Utilize YARN diagnostic and Spark error messages
- Fail YARN application based on last query's execution status
- Added an error classifier/ troubleshooting information to Dr. Elephant
- Clients retrieve troubleshooting information from Dr. Elephant APIs

STATEMENT 1 OUT OF 1 SHOW QUERY

Spark Tracking Url: http://monarch-adhoc-003-20200303-resource-manager-1:8088/proxy/application_1607358212046_636588/
Driver Logs Url: http://monarch-adhoc-003-20200303-data-worker-prod-0a026c30.ec2.pin220.com:8042/node/containerlogs/container_e20_1607358212046_636588_01_000001/sjaveria

YARN Diagnostics:
Failure Category: BROADCAST_TIMEOUT
Troubleshooting Info:
This error means that the application timed-out while waiting for broadcast join completions in the query or the driver/executor ran into GC.
For SparkSQL apps, try increasing "spark.sql.broadcastTimeout" config and retry.
If driver/executor ran into GC, increase the driver/executor memory using "spark.driver.memory"/"spark.executor.memory" configs.
TIP: You can find the default value on the SparkUI (under "environment" tab). Increase it by 25% to start with.
NOTE: Dr. Elephant might take some time to get ready for large applications.
Dr. Elephant exception analysis: https://drelephant.pinadmin.com/#/workflow-exceptions?applicationId=application_1607358212046_636588

Run on Apr 8, 4:07pm for 6m 24s by Sanchay Javeria.

Resource Tuning Suggestions

- User applications are often too generous with memory
- Provide visibility into over/ under memory consumption

1 Send to worker 2 Connect to engine 3 Running Query 1/1 4 Finish

STATEMENT 1 OUT OF 1 ▾

STATUS: RUNNING

SHOW QUERY × CANCEL QUERY

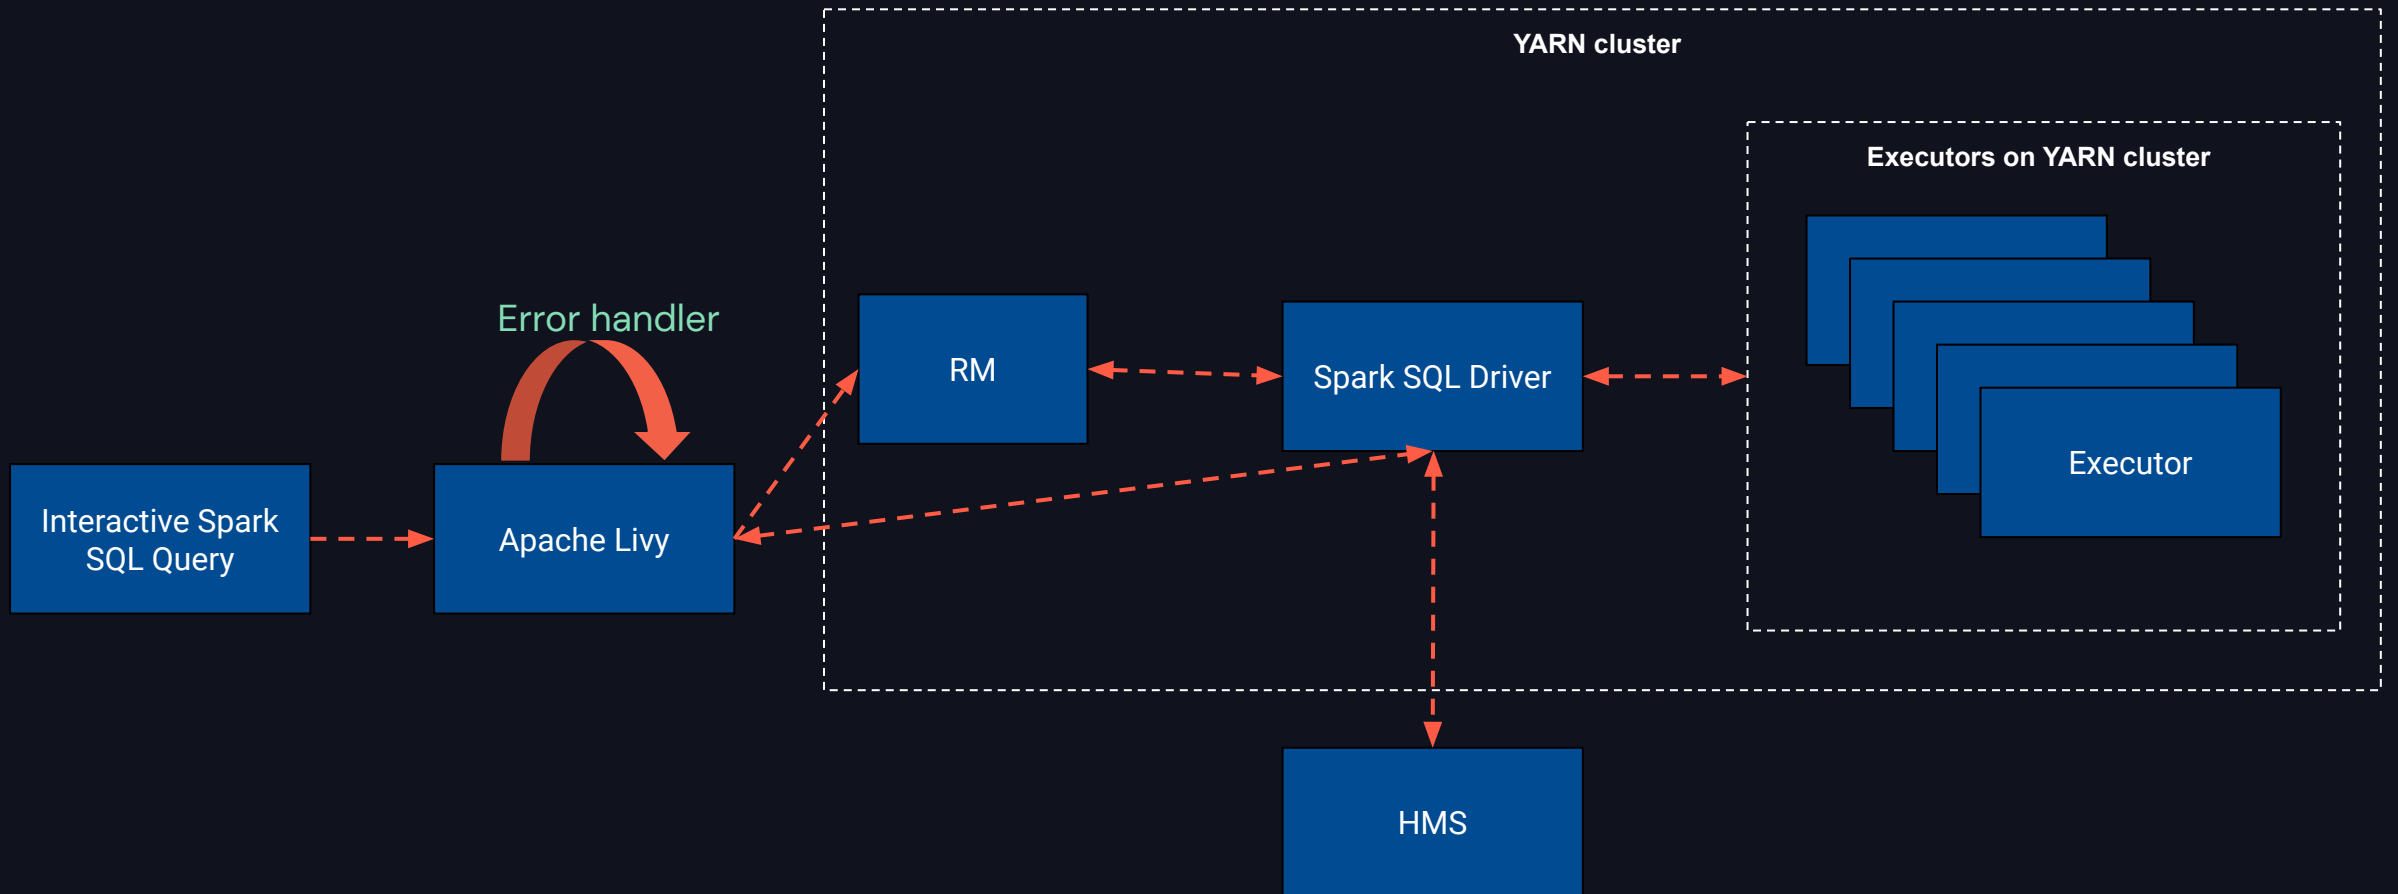
Spark UI
Driver Logs

(driver) memory used	max.	min.	avg.	feedback
3g	32.96%	32.96%	32.96%	

(executor) memory used	max.	min.	avg.	feedback
795m	6.72%	6.72%	6.72%	

0%

Auto Error Handling and Retries



Interactive Querying

Requirements Check

- Seamless query submission
- Fast metadata queries
- Fail fast
- Easy access to results
- Error handling and tuning suggestions
- Auto error handling and retries

Interactive UDFs

- Allow users to author UDFs along with queries.

Insert User Defined Function ✕

UDF Docs:
Spark UDF Guide

Function Name *

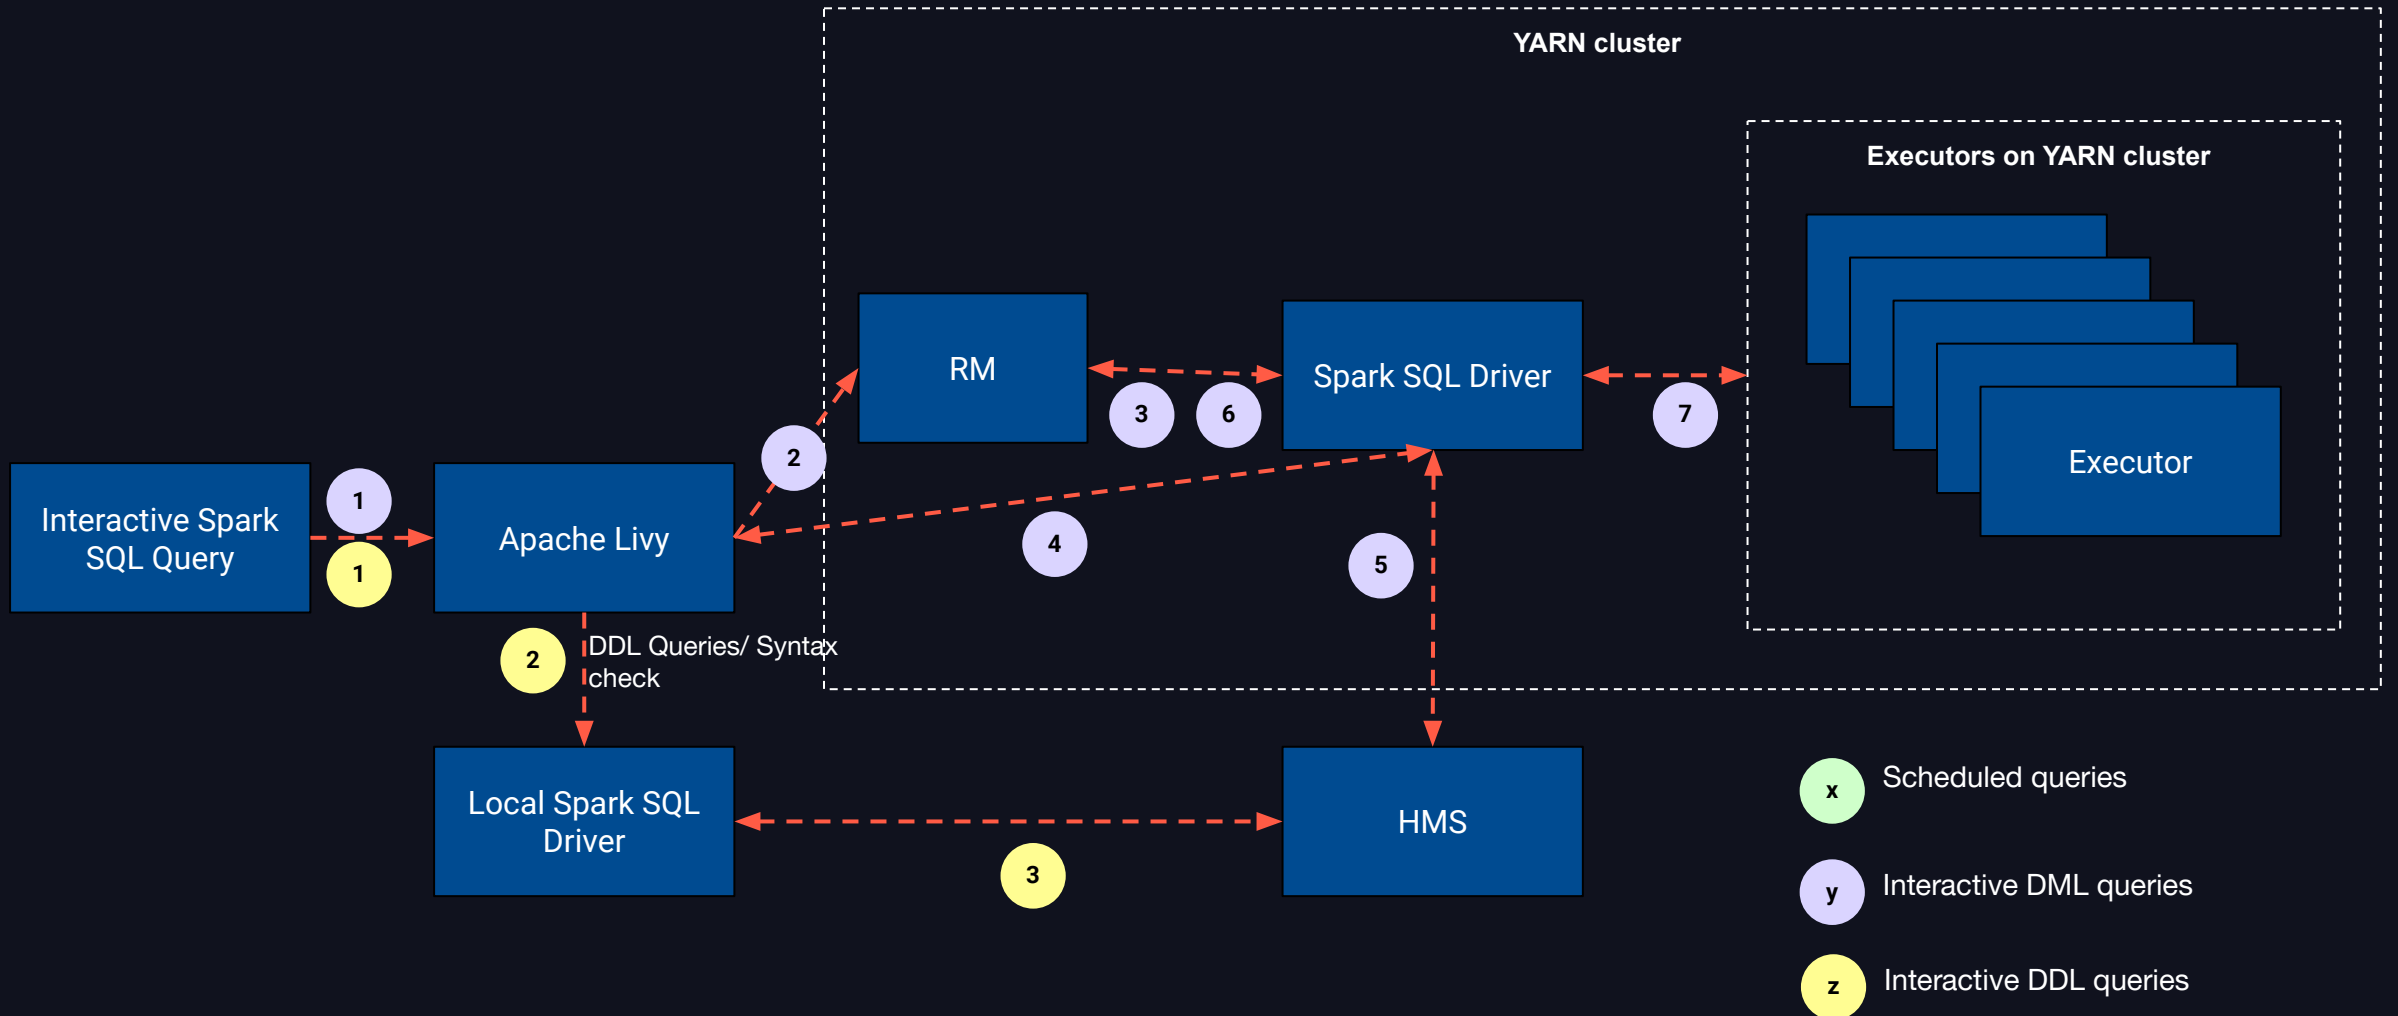
Language *

Code *

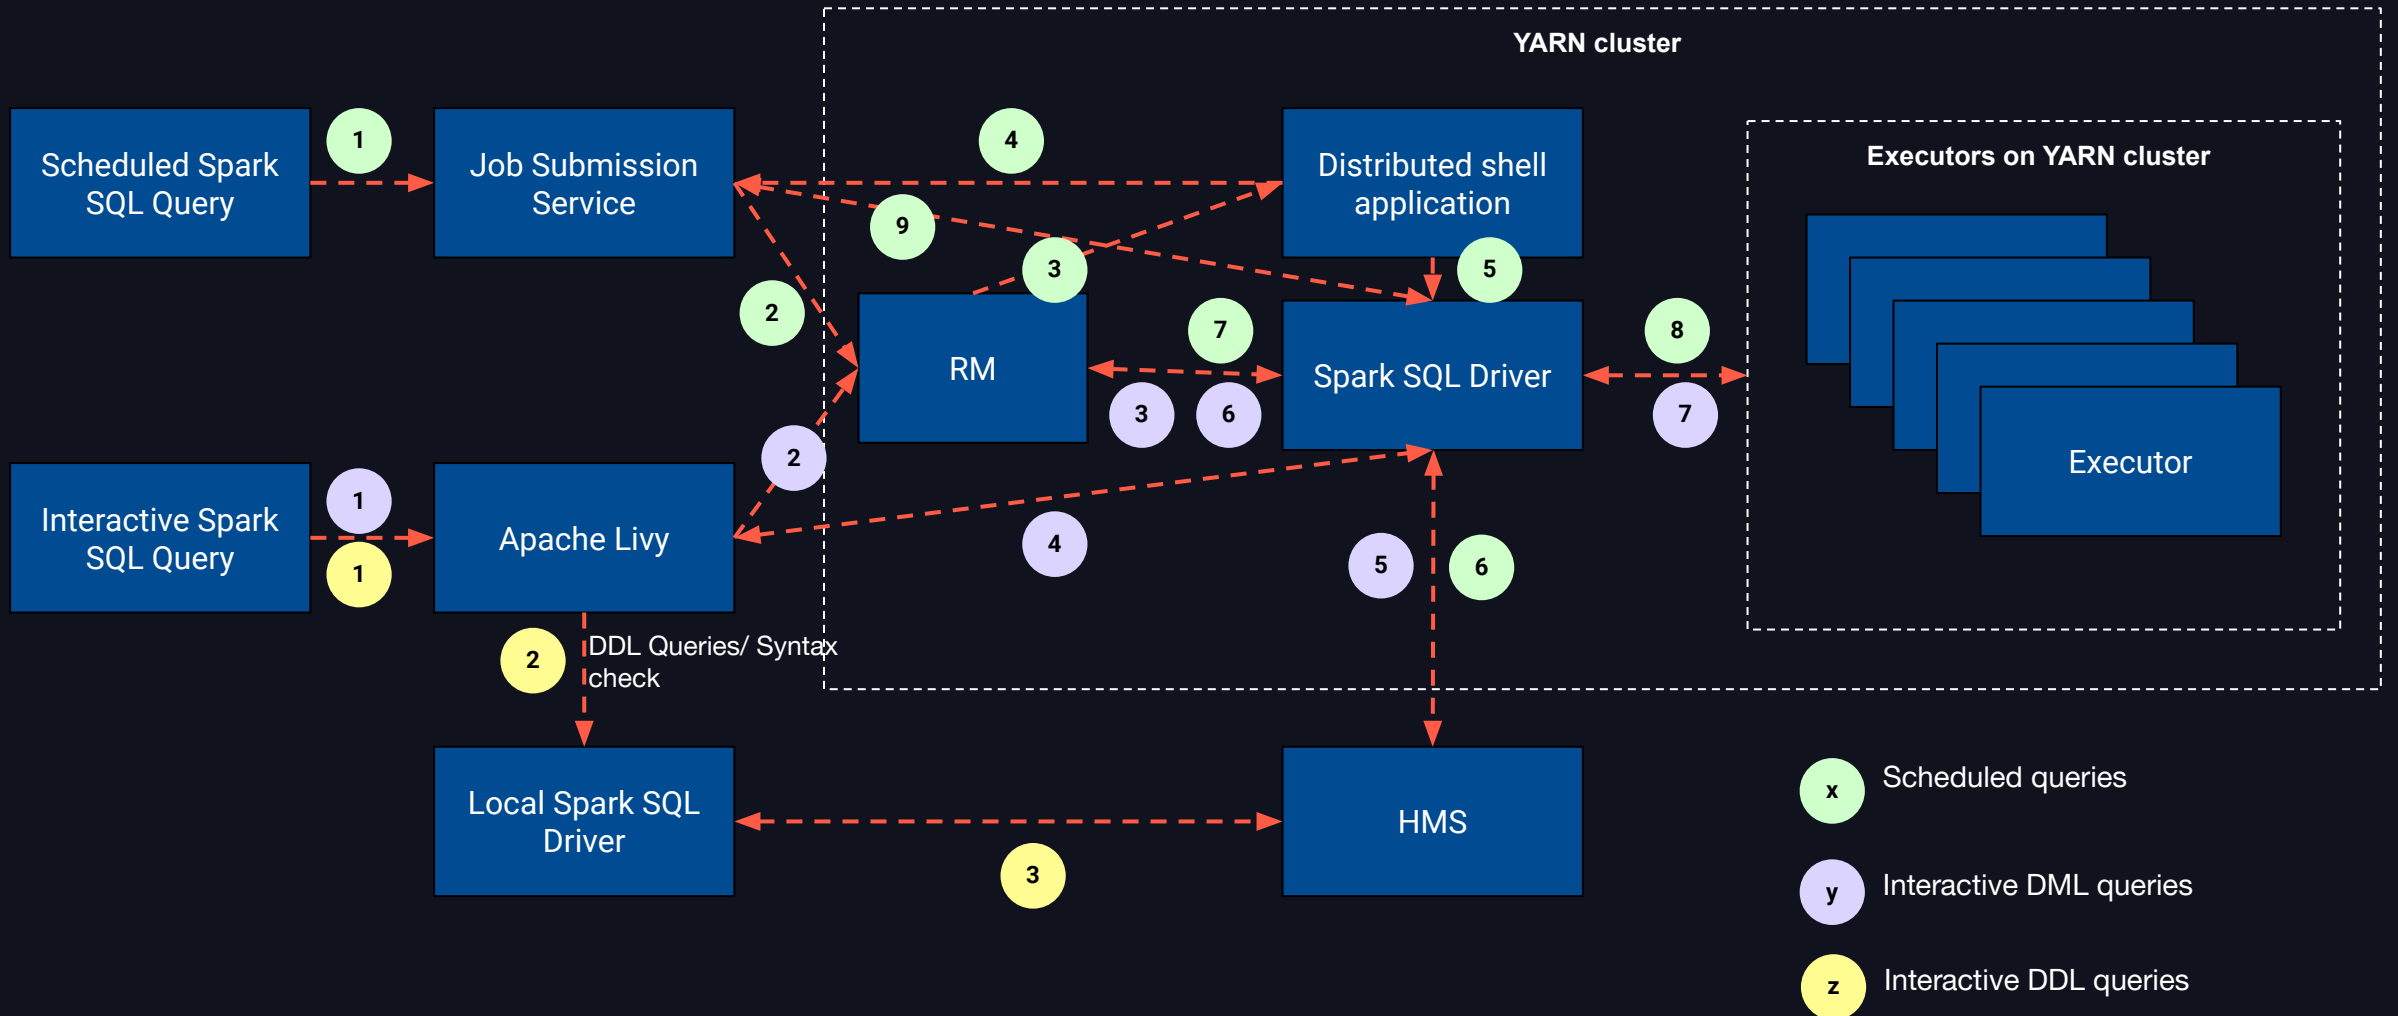
```
1 var FUNCTION_NAME_PLACEHOLDER = (x: Int) => {  
2   return x  
3 }
```

Querying with Spark SQL

Interactive Querying With Spark SQL



Querying With Spark SQL



DATA+AI
SUMMIT 2022

Thank you



Ashish Singh

Tech Lead, Big Data Query Platform,
Pinterest