



# Apache Kyuubi

——Flink SQL Engine设计&SQL计算中间件集成

杨华, 李心恺

2022.03.12

# CONTENT

01 JDBC之于Flink的现状

02 Flink SQL Engine的设计与实现

03 Flink SQL Engine展望

04 What`s Kyuubi & Why Kyuubi

05 Linkis 集成 Kyuubi实践过程

06 Kyuubi在一站式平台使用场景



# JDBC之于Flink的现状

# JDBC之于Flink现状



ververica / **flink-sql-gateway** Public

<> Code Issues 33 Pull requests 11 Actions Projects Wiki Security Insights

master - 6 branches 3 tags

Go to file Add file Code -

godfreyhe Support Flink-1.13-SNAPSHOT ca41846 on Feb 5, 2021 42 commits

ververica / **flink-jdbc-driver** Public

<> Code Issues 7 Pull requests 4 Actions Projects Wiki Security Insights

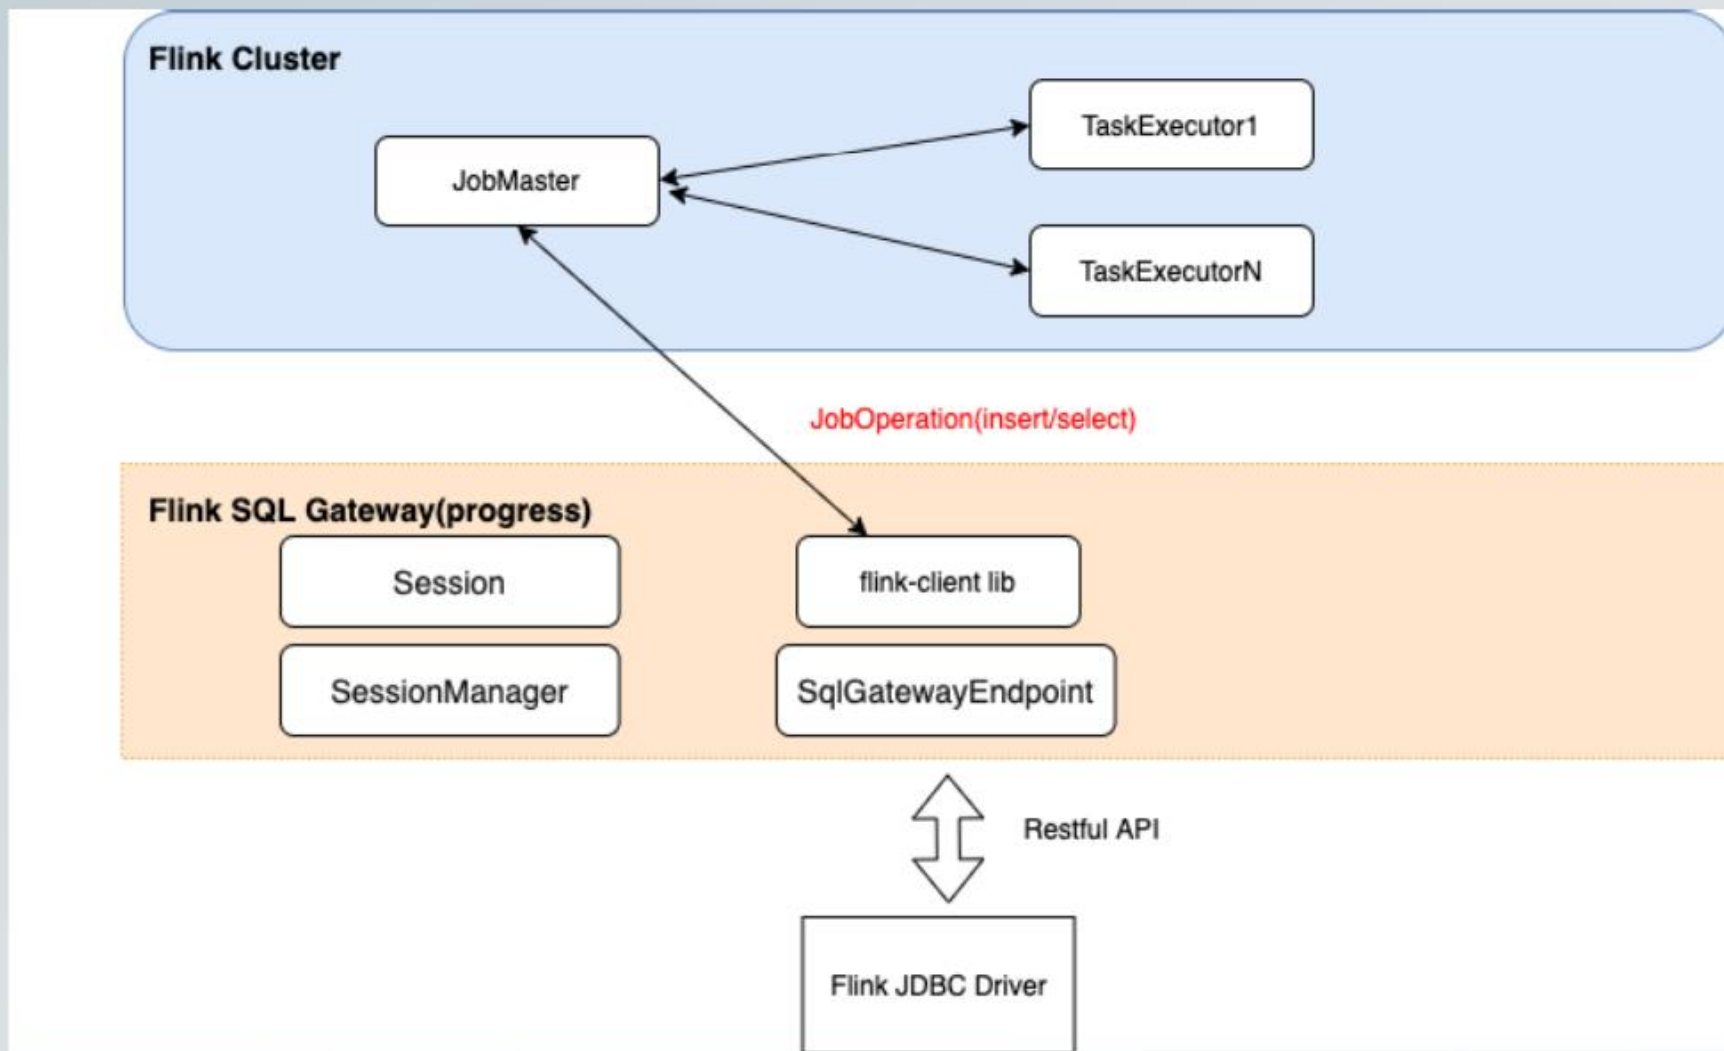
master - 4 branches 1 tag

Go to file Add file Code -

tsreaper and godfreyhe Set result fetching max sleep time according to query e... 1a423eb on Apr 8, 2020 4 commits



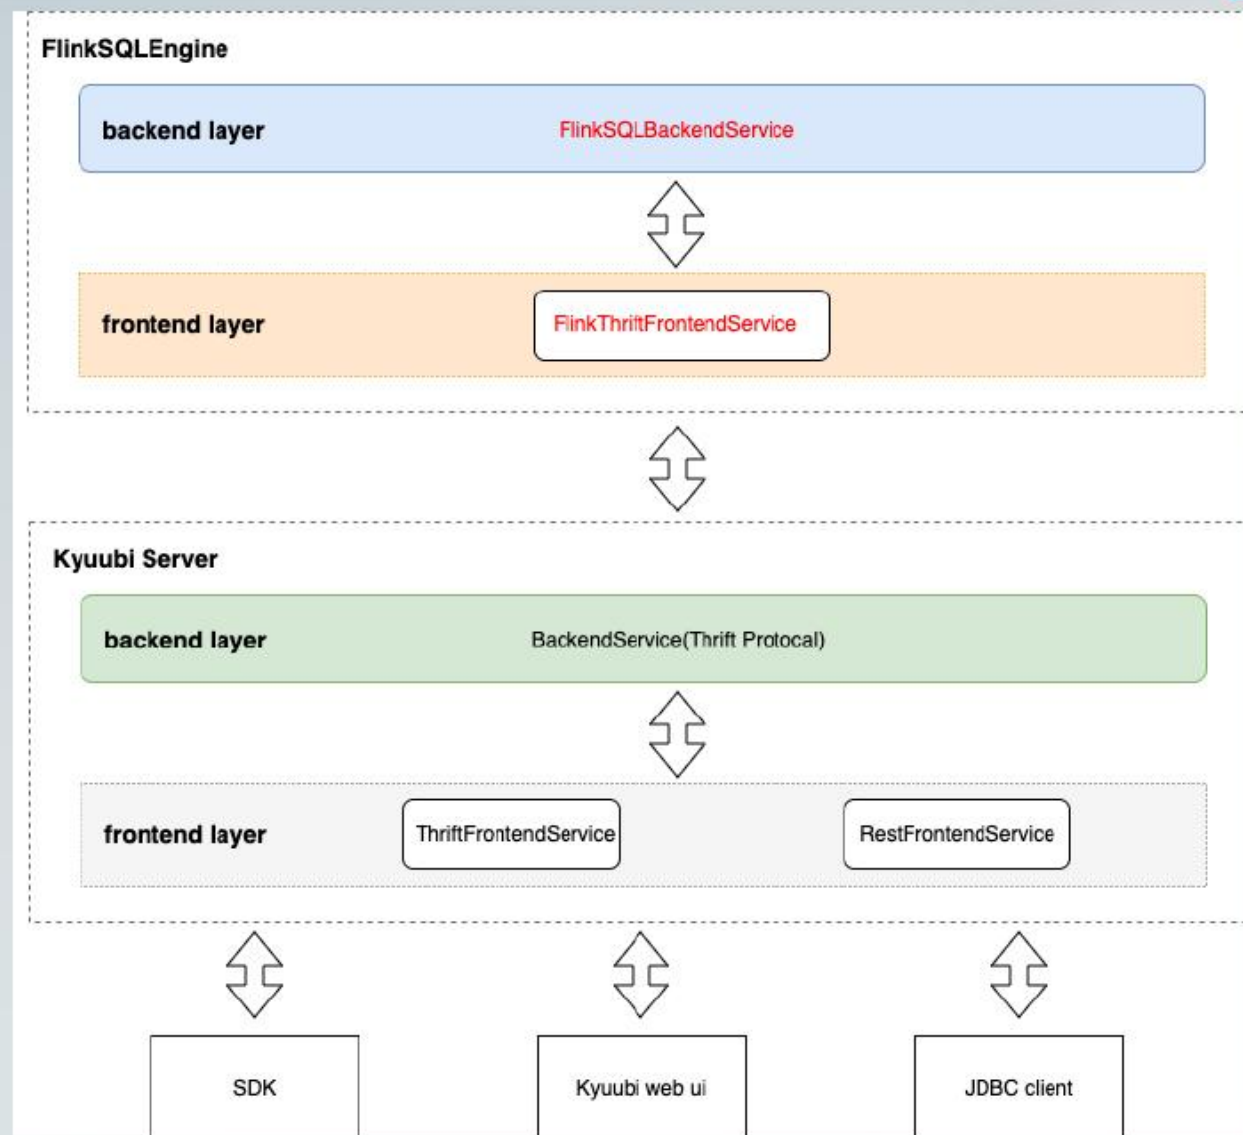
# Flink SQL Gateway & JDBC Driver 设计与实现





# Flink SQL Engine的设计与实现

# Flink SQL Engine









# Flink SQL Engine QuickStart Demo



```
./build/dist --name test --tgz
```

确保你将Hadoop classpath添加到PATH环境变量并启动Flink local cluster!



```
./bin/kyuubi start
```



```
./bin/beeline -u 'jdbc:hive2://192.168.1.9:10009/default?kyuubi.engine.type=FLINK_SQL' -n hive
```

# Flink SQL Engine QuickStart Demo



```
show catalogs;

show databases;

use default_database;

create table tbl_a (a int) with ('connector' = 'blackhole');
```

```
0: jdbc:hive2://192.168.1.9:10009/default> insert into tbl_a select 1;
23:32:54.863 INFO org.apache.kyuubi.operation.ExecuteStatement: Processing
4072-935e-0eff9eaadf9a]: INITIALIZED_STATE → PENDING_STATE, statement: ins
23:32:54.866 INFO org.apache.kyuubi.operation.ExecuteStatement: Processing
4072-935e-0eff9eaadf9a]: PENDING_STATE → RUNNING_STATE, statement: insert
23:32:59.875 INFO org.apache.kyuubi.operation.ExecuteStatement: Query[7a346
0eff9eaadf9a] in RUNNING_STATE
23:33:00.653 INFO org.apache.kyuubi.operation.ExecuteStatement: Query[7a346
0eff9eaadf9a] in FINISHED_STATE
23:33:00.653 INFO org.apache.kyuubi.operation.ExecuteStatement: Processing
4072-935e-0eff9eaadf9a]: RUNNING_STATE → FINISHED_STATE, statement: insert
taken: 5.787 seconds

+-----+
| default_catalog.default_database.tbl_a |
+-----+
| -1 |
+-----+
1 row selected (5.806 seconds)
```

insert-into\_default\_catalog.default\_database.tbl\_a **FINISHED** 1

ID: 9c7d97e1887ce7c255f788e782da5cf4 Start Time: 2022-03-10 23:32:57 End Time: 2022-03-10 23:32:58 Duration: 997ms

[Overview](#) [Exceptions](#) [TimeLine](#) [Checkpoints](#) [Configuration](#)

Source: Values[tuples=[[ 0 ]]] -> Calc[select=[1 AS EXP#0]] -> Sink[Sink[table=[default\_catalog.default\_database.tbl\_a], fields=[EXPR#0]]]

Parallelism: 1

Backpressure (min): N/A  
Busy (max): N/A

Name	Status	Bytes Received	Records Received	Bytes Sent	Records Sent	Parallelism	Start Time	Tasks
Source: Values[tuples=[[ 0 ]]] -> Calc[select=[1 AS EXP#0]] -> Sink[Sink[table=[default_catalog.default_database.tbl_a], fields=[EXPR#0]]]	FINISHED	0 B	0	0 B	0	1	2022-03-10 23:32:57	1

# Kyuubi Flink engine 目前进展



## [Umbrella][KPIP-2] Support FlinkSQL Engine #1322

Open 42 of 48 tasks yanghua opened this issue on Nov 2, 2021 · 13 comments



yanghua commented on Nov 2, 2021 · edited by yaoqinn

Member

### Code of Conduct

I agree to follow this project's Code of Conduct

### Search before asking

I have searched in the issues and found no similar issues.

### Describe the proposal

The detailed proposal is here. [\[KPIP-2\] Kyuubi Support Flink SQL Engine Design #1132](#)

### Task list

- [SUB-TASK][KPIP-2] Introduce EngineType to distinguish multiple engines #1323 @yanghua
- [SUB-TASK][KPIP-2] Support download Flink binary package in kyuubi-download module #1329 @yanghua
- [SUB-TASK][KPIP-2] Refactor ProcBuilder creation based on EngineType #1340 @yanghua
- [SUB-TASK][KPIP-2] Refactor SparkThriftBinaryFrontendService to extract connectionUrl into super class #1386 @yanghua
- [SUB-TASK][KPIP-2] Initialize kyuubi flink sql engine module #1492 @yanghua
- [SUB-TASK][KPIP-2] Introduce FlinkThriftBinaryFrontendService to do the further initialization #1503 @yanghua
- [SUB-TASK][KPIP-2] Implement GetCatalogs operation and do the further initialization #1525 @yanghua
- [SUB-TASK][KPIP-2] Move FetchIterator into kyuubi-common for multi-engine support #1542 @yanghua
- [SUB-TASK][KPIP-2] Bump flink version to 1.14.0 #1555 @yanghua
- [SUB-TASK][KPIP-2] Introduce flink-sql-client dependency to reduce some duplicated code #1576 @yanghua
- [SUB-TASK][KPIP-2] Implement basic ability of executing statement #1579 @yanghua
- [SUB-TASK][KPIP-2] Bump Flink 1.14.2 #1587 @pan3793
- Flink backend implementation #1629 @pan3793
- [SUB-TASK][KPIP-2] Implement GetFunctions operation #1643 @SteNicholas
- [SUB-TASK][KPIP-2] Implement GetSchemas operation #1645 @SteNicholas

Flink社区大佬们的鼎力支持:

- 蒋晓峰-阿里巴巴Blink团队(23个PR)
- 林小柏-网易游戏实时团队(11个PR)

当前提供的功能:

- DDL/DML
- JDBC: GetXXX/ShowXXX
- set/show/reset properties
- UDF
- Flink Session Standalone/on YARN

<https://github.com/apache/incubator-kyuubi/issues/1322>



# Flink SQL Engine展望



# Flink 未来的 deploy mode



## Application Mode



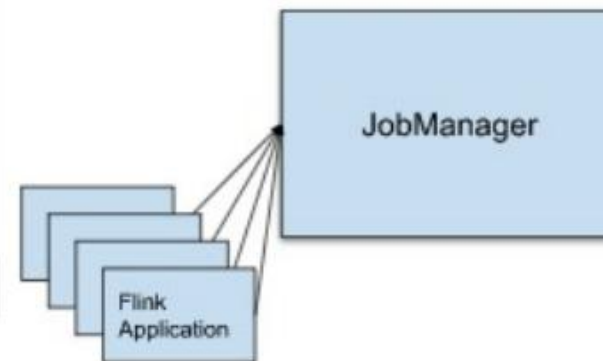
A dedicated JobManager is started for submitting the job. The JobManager will only execute this job, then exit.  
The Flink Application runs on the JobManager.

## Per-Job Mode



A dedicated JobManager is started for submitting the job. The JobManager will only execute this job, then exit.  
The Flink Application runs on the client submitting the per-job cluster

## Session Mode



Multiple jobs share one JobManager.

<https://issues.apache.org/jira/browse/FLINK-26000>

# Flink SQL Engine 接下来的规划



- Flink SQL Engine on YARN(application)
- Flink SQL Engine on K8S
- Flink SQL Engine shared level enhancement
- Support session jars management
- .....

<https://github.com/apache/incubator-kyuubi/issues/2100>





# **What`s Kyuubi &Why Kyuubi**

# T3出行 >>

是一家基于车联网驱动的智慧出行平台，拥有海量且丰富的数据源。因为车联网数据多样性，T3出行构建了以Apache Hudi为基础的企业级的数据湖，并在此之上构建了BI分析平台，任务调度，机器学习平台，数据质量，权限管理等一系列平台，为业务提供强有力的支撑。

随着平台越来越多，业务结构越来越复杂，年初经过一段时间的一站式数据应用交互管理平台的技术调研选型，我们最终确定以微众银行开源的DataSphere Studio作为Insight的一站式数据应用交互管理平台，并根据公司业务需要进行一些定制开发。



# DSS引入Kyububi之前架构







## 01 跨存储

数据分布在OBS、Hudi、Clickhouse等不同存储，需要写代码关联分析增加数据处理门槛和成本，Linkis对此解决有限

## 02 SQL不统一

Hive不支持upsert update delete 等语法操作hudi表。同时Mongodb、Clickhouse等语法又各不相同，开发转换成本较高

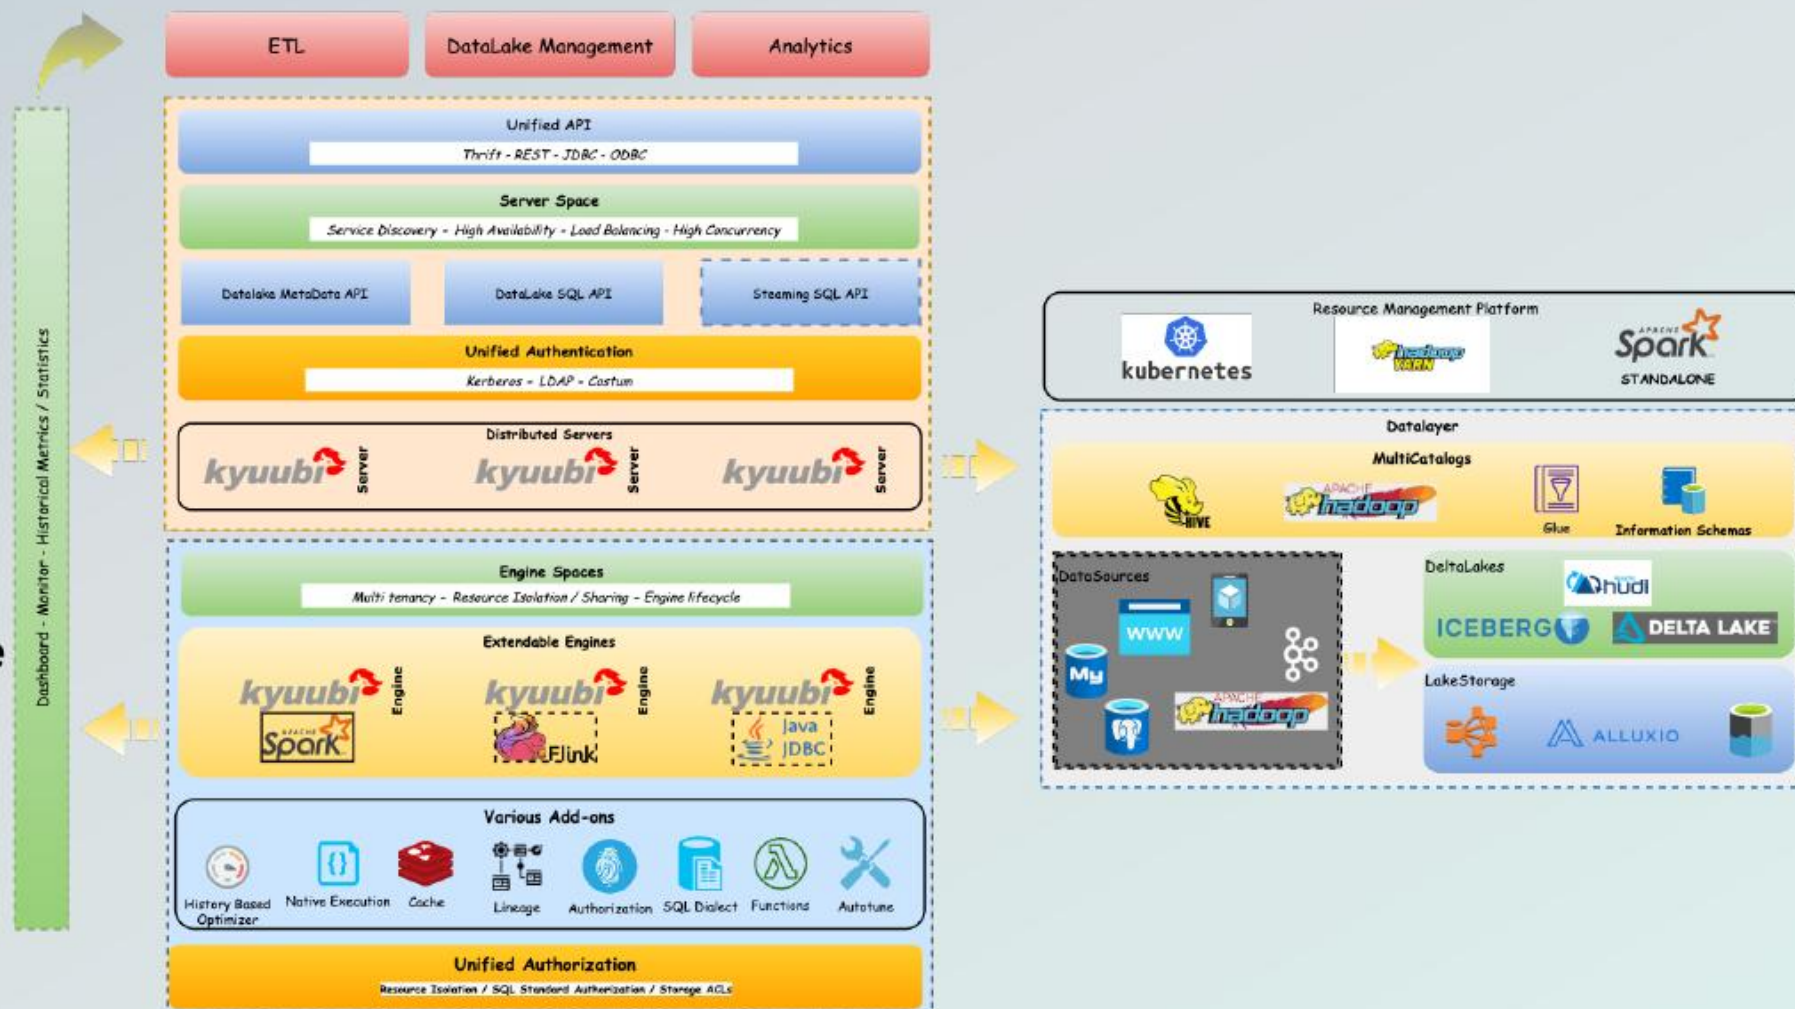
## 03 版本耦合

Linkis 和Hive版本、Spark版本强耦合，集成spark3.X难度大，同时Linkis中的Spark引擎对于cluster运行模式、AQE、动态资源等特性的支持还不完善

# 什么是Apache Kyuubi



- Thrift JDBC/ODBC 服务
- 可扩展
- 多租户和分布式
- Servless SQL on Lakehouse



# 选择Kyuubi 优势



	HiveServer2 (Hive on Spark)	Linkis	Kyuubi
接口	HiveJDBC	Http	HiveJDBC
SQL语法	Hive SQL	Spark SQL、Hql	Spark SQL
SQL解析	Server端	Engine端	Engine端
任务提交	拆分成多个RemoteDriver提交	共享Server这个Spark程序, 分布式线程调度	USER、GROUP、CONNECTION级别隔离
Spark兼容性	特定单一版本支持	和Spark、Hive特定版本绑定	多版本适配
Catalog管理	HMS	HMS	HMS、第三方Catalog
高可用	有	有	有
多租户	有	有	有
权限控制	SQL Standard,细粒度	Hive 自身权限控制	SQL Standard.细粒度
性能	一般	一般	好
客户端并发	水平扩展	水平扩展	水平扩展
动态资源配置	SQL粒度	无	Engine粒度
计算资源管理	YARN	自我管理.Fair Scheduler Pools	YARN、Kubernetes等
资源占用周期	SQL执行期间	引擎启动期间	资源通过Engine维度来 申请和释放





# Linkis 集成 Kyuubi 实践过程

# Linkis 集成Kyuubi流程



Linkis 支持新增一个自定义引擎。Linkis 支持多种引擎类型。

其中ComputationExecutor: 是常用的交互式引擎Executor, 处理交互式执行任务, 并且具备状态查询、任务kill等交互式能力。

Kyuubi引擎是一个交互式引擎, 因此在实现Executor时, 继承了ComputationExecutor, 并做了以下maven依赖的引入。

```
<dependency>
<groupId>org.apache.linkis</groupId>
<artifactId>linkis-computation-engineconn</artifactId>
<version>${linkis.version}</version>
</dependency>
```

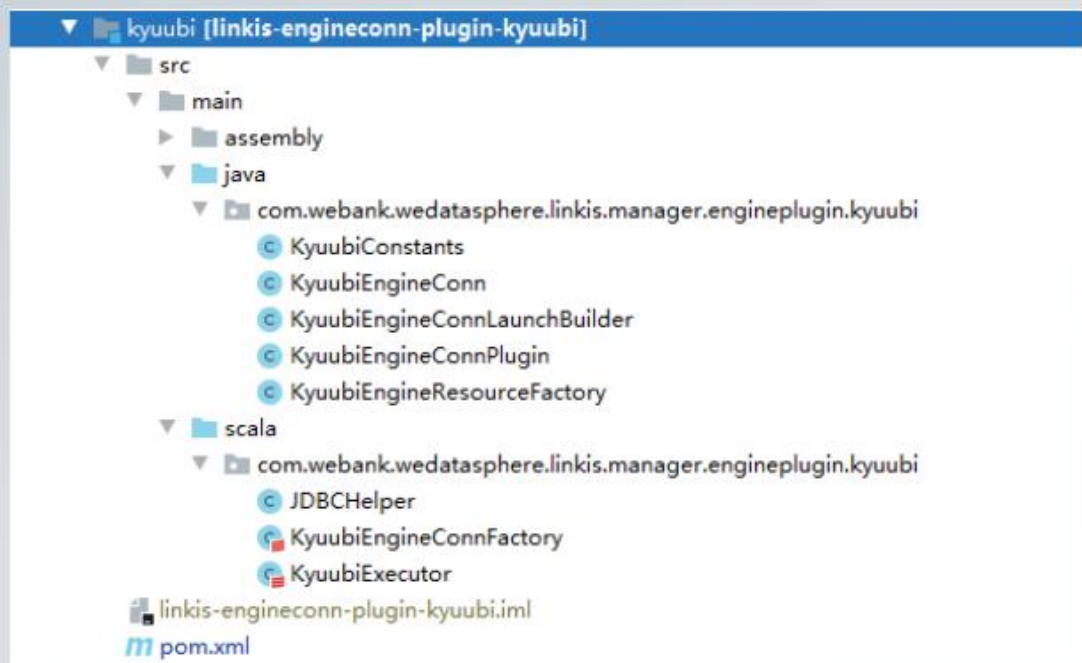
具体可参考linkis官方文档《如何快速实现新的底层引擎》

# Linkis 集成Kyuubi流程

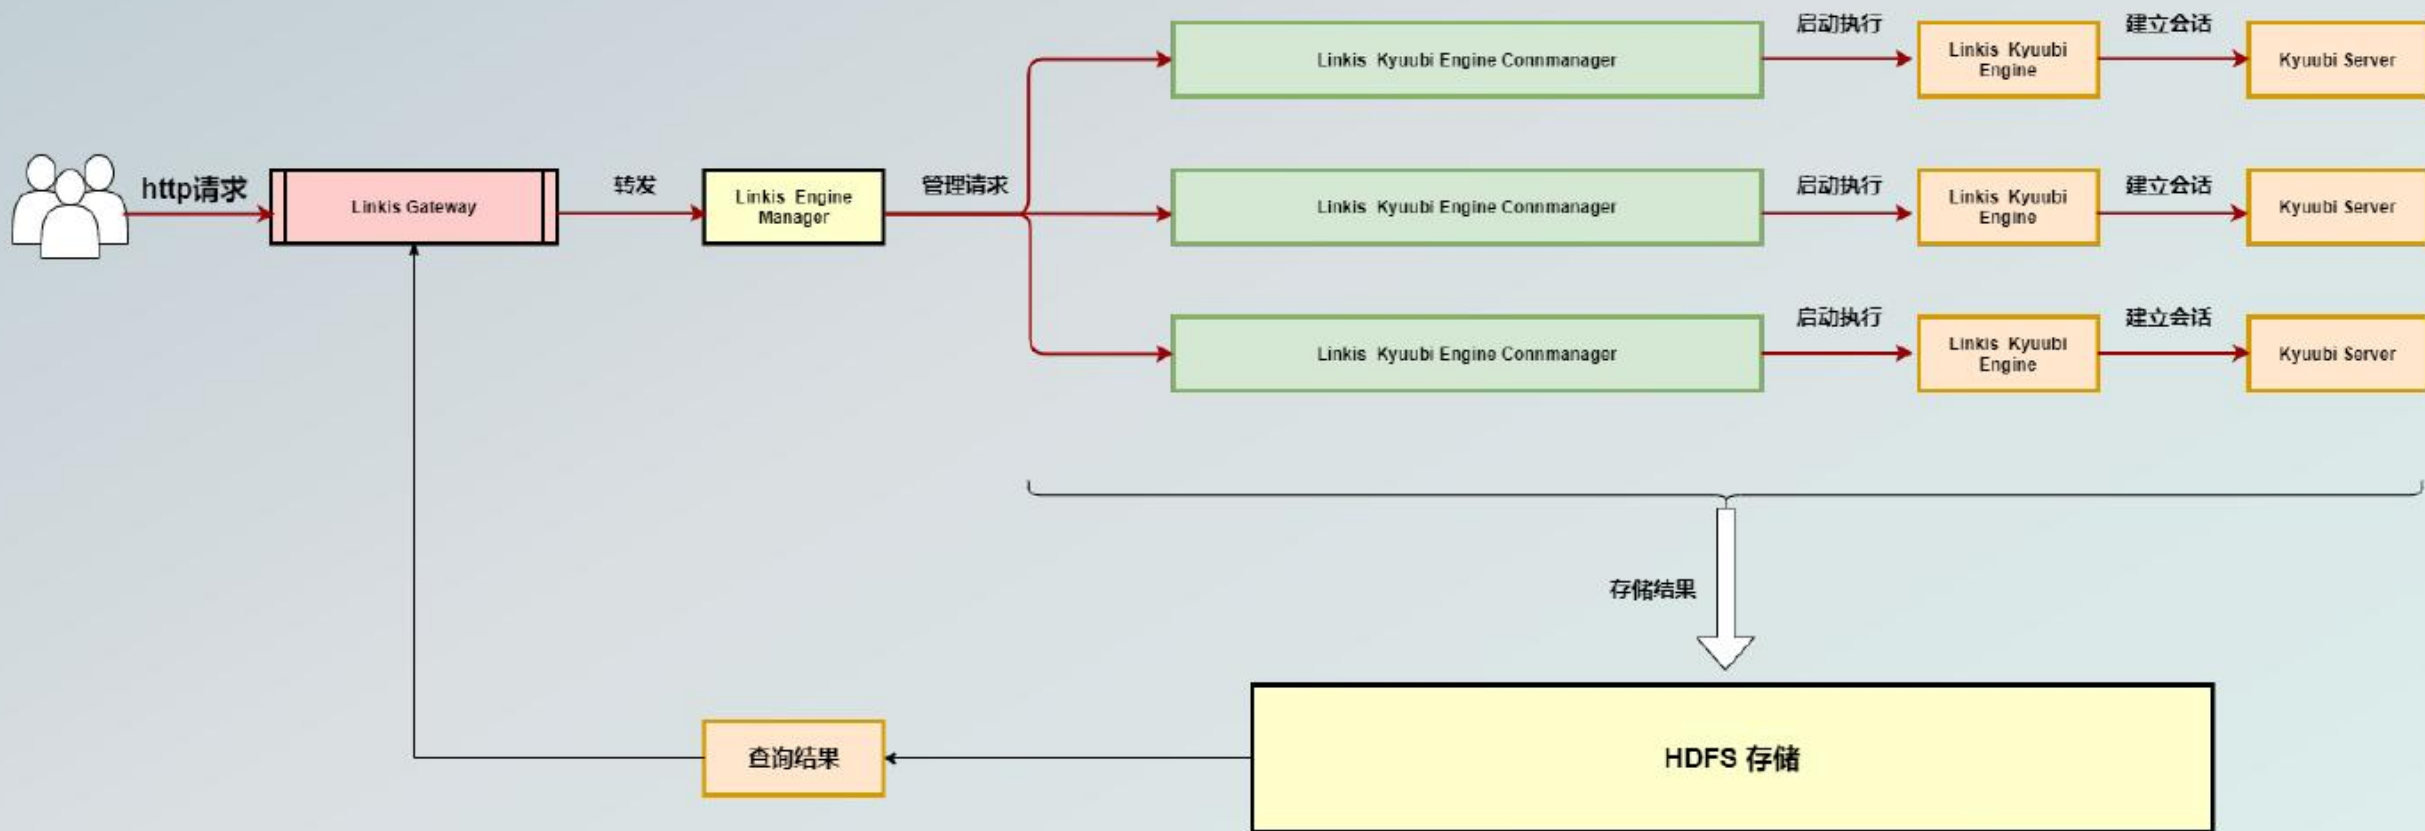


主要是要实现：

- ✓ KyuubEngineConnPlugin, 启动EngineConn时的入口
- ✓ KyuubEngineConnFactory, 实现如何启动一个引擎连接器, 和如何启动一个引擎执行器的逻辑
- ✓ KyuubEngineLaunchBuilder, 用于封装EngineConnManager可以解析成启动命令
- ✓ KyuubiExecutor为执行器, 作为真正的计算场景执行器, 是实际的计算逻辑执行单元



# Kyuubi Engine执行逻辑





# Kyuubi+ Linkis 的一站式架构





# Kyuubi在一站式平台使用场景



# DSS集成Kyuubi



DSS数据开发模块，直接对接Kyuubi服务，实现一站式开发和CI\CD

The screenshot displays the DataSphere Studio v1.0.0 interface. The top navigation bar includes '首页', 'Scriptis', and '管理台'. The left sidebar shows a file tree with 'hive' > 'test' > 'testJDBCjdbc' selected. The main editor area contains a SQL script with the following content:

```
1 select
2   ds
3   ,city_name
4   ,sum(call_cnt) a
5   ,sum(receive_cnt) b
6
7   ,sum(receive_cnt)/sum(call_cnt) c
8 from test...
9 where ds>=replace('2022-02-12','-', '') and ds<=replace('2022-02-12','-', '')
10 and city_name in (...)
11 and route_source in (...)
12 group by city_name,ds
13
```

# 任务编排发布



任务编排集成kyuubi, 可以直接拉去kyuubi组件编排任务流, 关联编辑SQL脚本

The screenshot displays a workflow editor interface. At the top, the breadcrumb path is 'bdapWorkspace / temp'. Below this, there are tabs for '开发中心' and 'temp1'. The main workspace shows a workflow diagram with two tasks: 'sql\_5700' and 'kyuubi'. The 'sql\_5700' task is connected to the 'kyuubi' task. On the left side, there is a sidebar with various components categorized into '数据开发' (Data Development), '数据治理' (Data Governance), '数据集成' (Data Integration), and '数据输出' (Data Output). The '数据开发' category includes icons for 'sql', 'python', 'spark', 'scala', 'hql', and 'shell'. The '数据治理' category includes 'monitor...', 'metrics...', and 'dashbor...'. The '数据集成' category includes 'connector' and 'subflow'. The '数据输出' category includes 'rendered'. At the bottom right, there is a zoom control showing '100%'.

# DS调度集成



 任务编排完成的任务，发布对接到DS，DS以SQL模块执行脚本



The screenshot displays the DolphinScheduler web interface. On the left, a sidebar contains navigation options: 项目首页, 工作流, 工作流定义 (highlighted), 工作流实例, and 任务实例. The main workspace shows a task node being configured. A blue arrow points from the '开始' (Start) button in the task editor to the 'sql参数' field in the configuration panel.

**当前节点设置**

- 节点名称: task\_verifydata\_dws\_dri\_sal\_index\_ds\_完单司机数\_业务线维度
- 运行标志:  正常  禁止执行
- 描述: 业务线维度, 按司机数维度, 完单司机数, 维度维度, 完单司机数, 完单司机数
- 任务优先级: MEDIUM
- Worker分组: offline
- 失败重试次数: 2 (次)
- 失败重试间隔: 2 (分)
- 超时告警:
- 失败告警:
- 数据源: HIVE, Kyuubi
- sql类型: 非查询
- sql参数: 请输入格式为 key1=value1:key2=value2...

底部操作按钮: 取消, 确认添加

# 监控管理



因为开发和调度大量使用的kyuubi，加强了kyuubi后台管理，通过管理ui对用户操作统一管理监控





# 监控管理



主页 服务管理 SparkContext **用户会话** SQL历史 租户管理 配置管理 日志级别

ApplicationId	用户	状态	Running SQL	Error SQL	Total SQL	开始时间	结束时间	Links
application_1635251429604_770960	bi_shaowei	stopped	0	0	1	2022-03-08 19:31:29	2022-03-08 19:32:09	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:31:29	2022-03-08 19:31:33	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:31:15	2022-03-08 19:31:19	Sql
application_1635251429604_770960	bi_shaowei	stopped	0	0	1	2022-03-08 19:31:01	2022-03-08 19:31:04	Sql
application_1635251429604_770790	yuansi	stopped	0	0	1	2022-03-08 19:30:39	2022-03-08 19:32:37	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:30:36	2022-03-08 19:31:04	Sql
application_1635251429604_770989	bi_huxiaok	stopped	0	0	1	2022-03-08 19:30:31	2022-03-08 19:32:49	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:29:43	2022-03-08 19:29:49	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:26:42	2022-03-08 19:26:48	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:25:58	2022-03-08 19:26:25	Sql
application_1635251429604_770957	sunlulin	stopped	0	0	1	2022-03-08 19:25:04	2022-03-08 19:25:50	Sql
application_1635251429604_770790	yuansi	stopped	0	0	1	2022-03-08 19:24:55	2022-03-08 19:29:27	Sql
application_1635251429604_770960	bi_shaowei	stopped	0	0	1	2022-03-08 19:24:06	2022-03-08 19:29:27	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:24:04	2022-03-08 19:24:08	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:23:27	2022-03-08 19:23:43	Sql
application_1635251429604_770959	bi_zhuj	stopped	0	0	2	2022-03-08 19:22:35	2022-03-08 19:22:55	Sql
application_1635251429604_770958	zhoumin	stopped	0	0	1	2022-03-08 19:22:13	2022-03-08 19:23:14	Sql
application_1635251429604_770957	sunlulin	stopped	0	0	1	2022-03-08 19:20:51	2022-03-08 19:21:43	Sql



通过Kyuubi Server Api接口，管理会话列表，并且可以手动关闭会话

# 监控管理



用户执行的sql可以看到详细信息，并可以管理取消用户操作

Kyuubi 主页 服务管理 SparkContext 用户会话 SQL历史 租户管理 配置管理 日志级别

ApplicationId:  用户:   
状态:  SQL:  开始时间 >=:   
开始时间 <:

用户	SQL	状态	耗时	开始时间	结束时间	Links
sunlulin	select f1.ds '日期分区', f1.city_name '城市', f1.agent_name '运营商', ... <a href="#">复制</a>	running <input type="button" value="取消"/>		2022-03-08 19:25:04		<a href="#">Session</a>
yuansi	with A as( select '20220307' ds union all select from_unixtime(unix_ti... <a href="#">复制</a>	error <input type="button" value="复制"/>	00:00:03	2022-03-08 19:24:55	2022-03-08 19:24:58	<a href="#">Session</a>
bi_shaowei	select t1.city_code, t1.city_name, sum(t1.w_onjob) as w_dri_cnt, sum(... <a href="#">复制</a>	error <input type="button" value="复制"/>	00:00:03	2022-03-08 19:24:24	2022-03-08 19:24:27	<a href="#">Session</a>
bi_zhuj	use datalake_t3_ml <input type="button" value="复制"/>	closed	00:00:00	2022-03-08 19:24:04	2022-03-08 19:24:04	<a href="#">Session</a>
bi_zhuj	select city_id, min(id), max(id), count(1) from datalake_t3_ml global_on... <a href="#">复制</a>	closed	00:00:04	2022-03-08 19:24:04	2022-03-08 19:24:08	<a href="#">Session</a>
bi_zhuj	use datalake_t3_ml <input type="button" value="复制"/>	closed	00:00:00	2022-03-08 19:23:27	2022-03-08 19:23:27	<a href="#">Session</a>
bi_zhuj	select city_id, min(id), max(id), count(1) from datalake_t3_ml global_on... <a href="#">复制</a>	closed	00:00:15	2022-03-08 19:23:27	2022-03-08 19:23:42	<a href="#">Session</a>
bi_zhuj	use datalake_t3_ml <input type="button" value="复制"/>	closed	00:00:00	2022-03-08 19:22:54	2022-03-08 19:22:54	<a href="#">Session</a>
bi_zhuj	select city_id, min(id), max(id), count(1) from datalake_t3_ml global_on... <a href="#">复制</a>	error <input type="button" value="复制"/>	00:00:01	2022-03-08 19:22:54	2022-03-08 19:22:55	<a href="#">Session</a>
zhoumin	select city_name, type_module, origin_lat, origin_lng, sum(callinum) c... <a href="#">复制</a>	closed	00:00:44	2022-03-08 19:22:30	2022-03-08 19:23:14	<a href="#">Session</a>
sunlulin	select f1.ds '日期分区', f1.city_name '城市', f1.agent_name '运营商', ... <a href="#">复制</a>	closed	00:00:33	2022-03-08 19:21:09	2022-03-08 19:21:42	<a href="#">Session</a>



# 总结



T3出行大数据平台引入Apache Kyuubi后，和Linkis功能互补，实现了代码开发、业务上线与调度系统的打通，同时可以收口做到大数据开发CI\CD管理，帮助业务部门低门槛上线大数据相关的需求，减轻了数据开发的压力，向一站式开发平台的目标更进一步。也期望Apache kyuubi 和Linkis 作为计算中间件的引领者越来越好！





**THANK YOU**  
QUESTIONS?

