

AI iSQL Plan Cache

阿里云数据库内核组

王雄

中文名：王雄

英文名：Benny Wang



2016年加入阿里云，担任数据库内核组高级专家职位



2013年任职于ORACLE MySQL optimizer team担任Principle工程师职位

TERADATA

2012年任职性能分析工程师

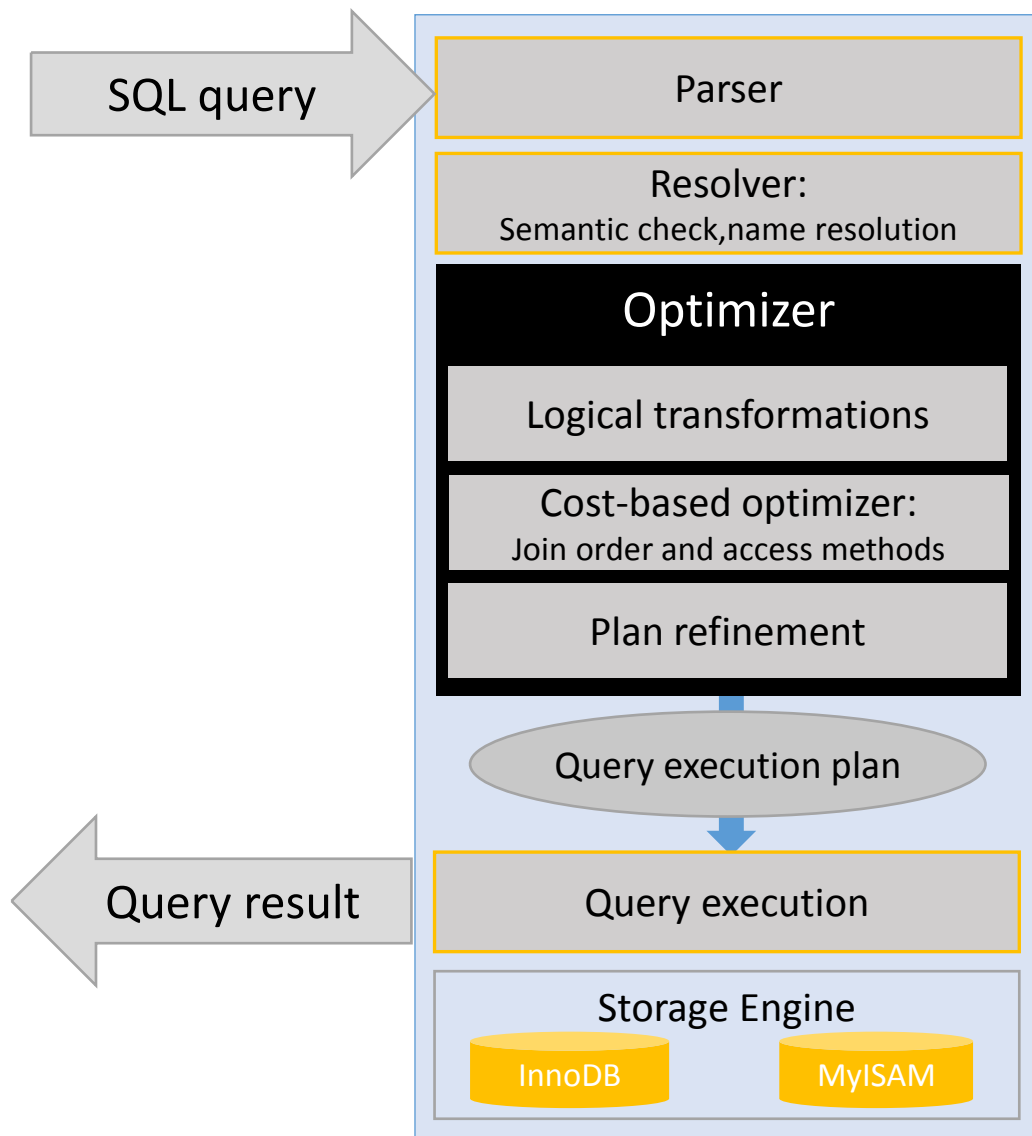


2009年任职人大金仓数据库高级研发工程师

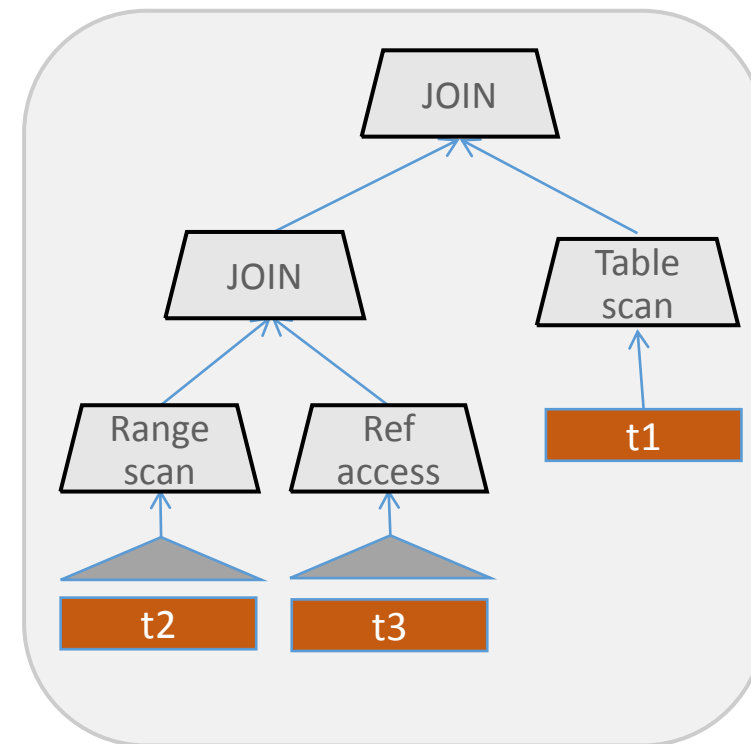
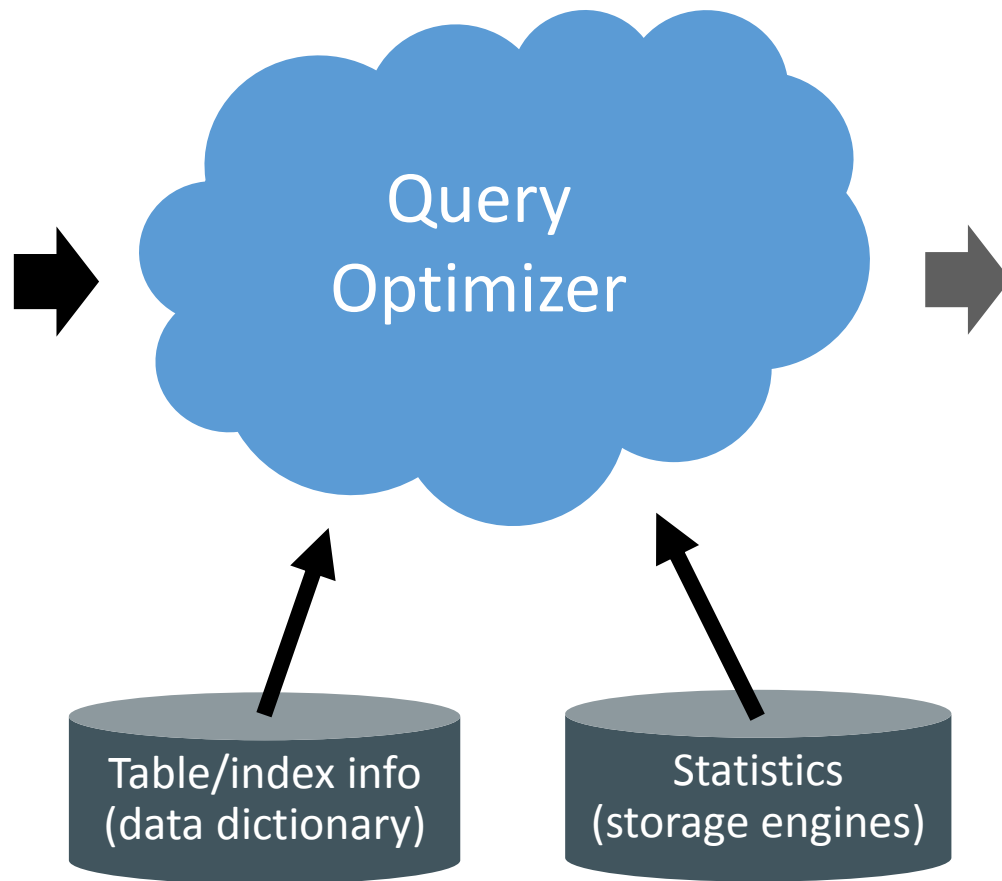


- 01** 优化器背景知识介绍
- 02** 什么是Plan cache
- 03** Plan cache的意义
- 04** AliSQL Plan cache的实现和使用
- 05** 性能测试
- 06** 未来之路



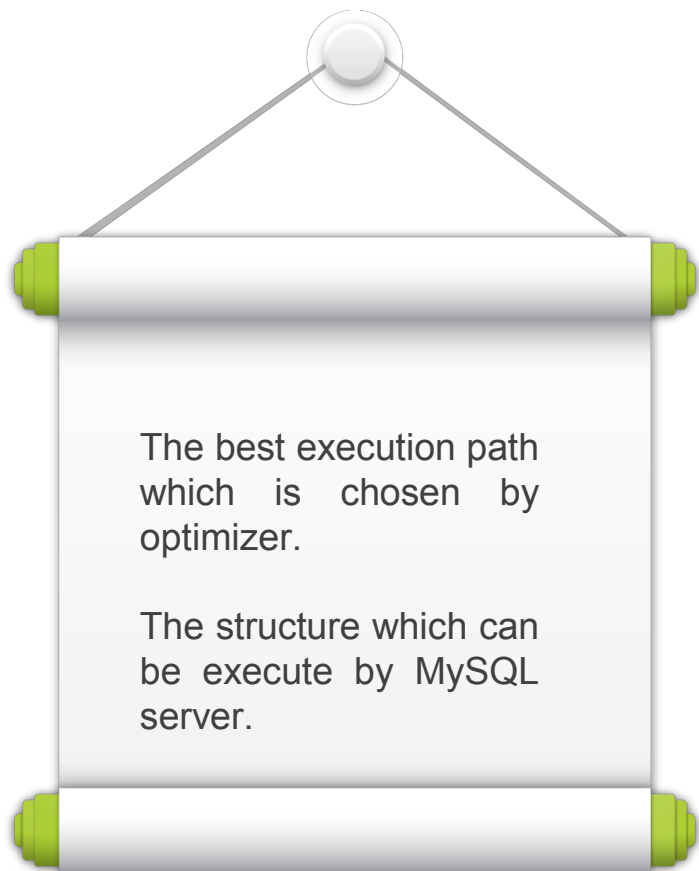


```
SELECT a, b  
FROM t1, t2, t3  
WHERE t1.a = t2.b  
AND t2.b = t3.c  
AND t2.d > 20  
AND t2.d < 30;
```

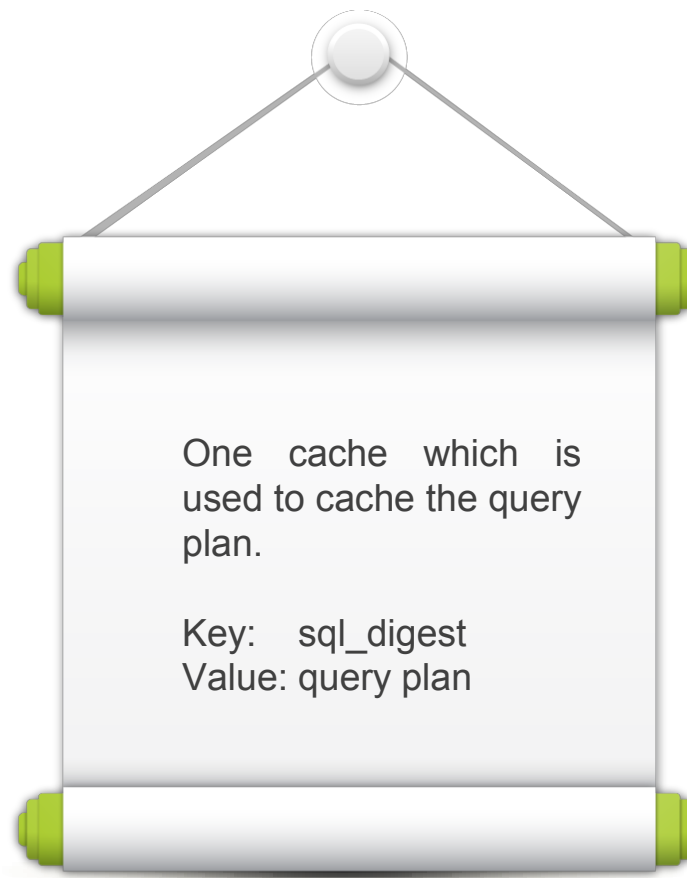


什么是 Plan Cache ?

Query Plan



Plan Cache



- `SELECT * FROM t1 WHERE a > 1;`
- `select * from t1 where a > 1;`
- `SELECT * FROM /* comments */ t1 where a > 1;`
- `select * from t1 where a > 3`


`SELECT * FROM `t1` WHERE `a` > ?`



- 传统数据库都标配的特性，ORACLE，DB2，SQL SERVER等
- 频繁执行某些查询语句
- 并发量大，CPU成为瓶颈
- 优化器在生成Plan的时候会消耗较长的时间
- 协助用户固化Query Plan
- 实现“真正的” Prepared statement



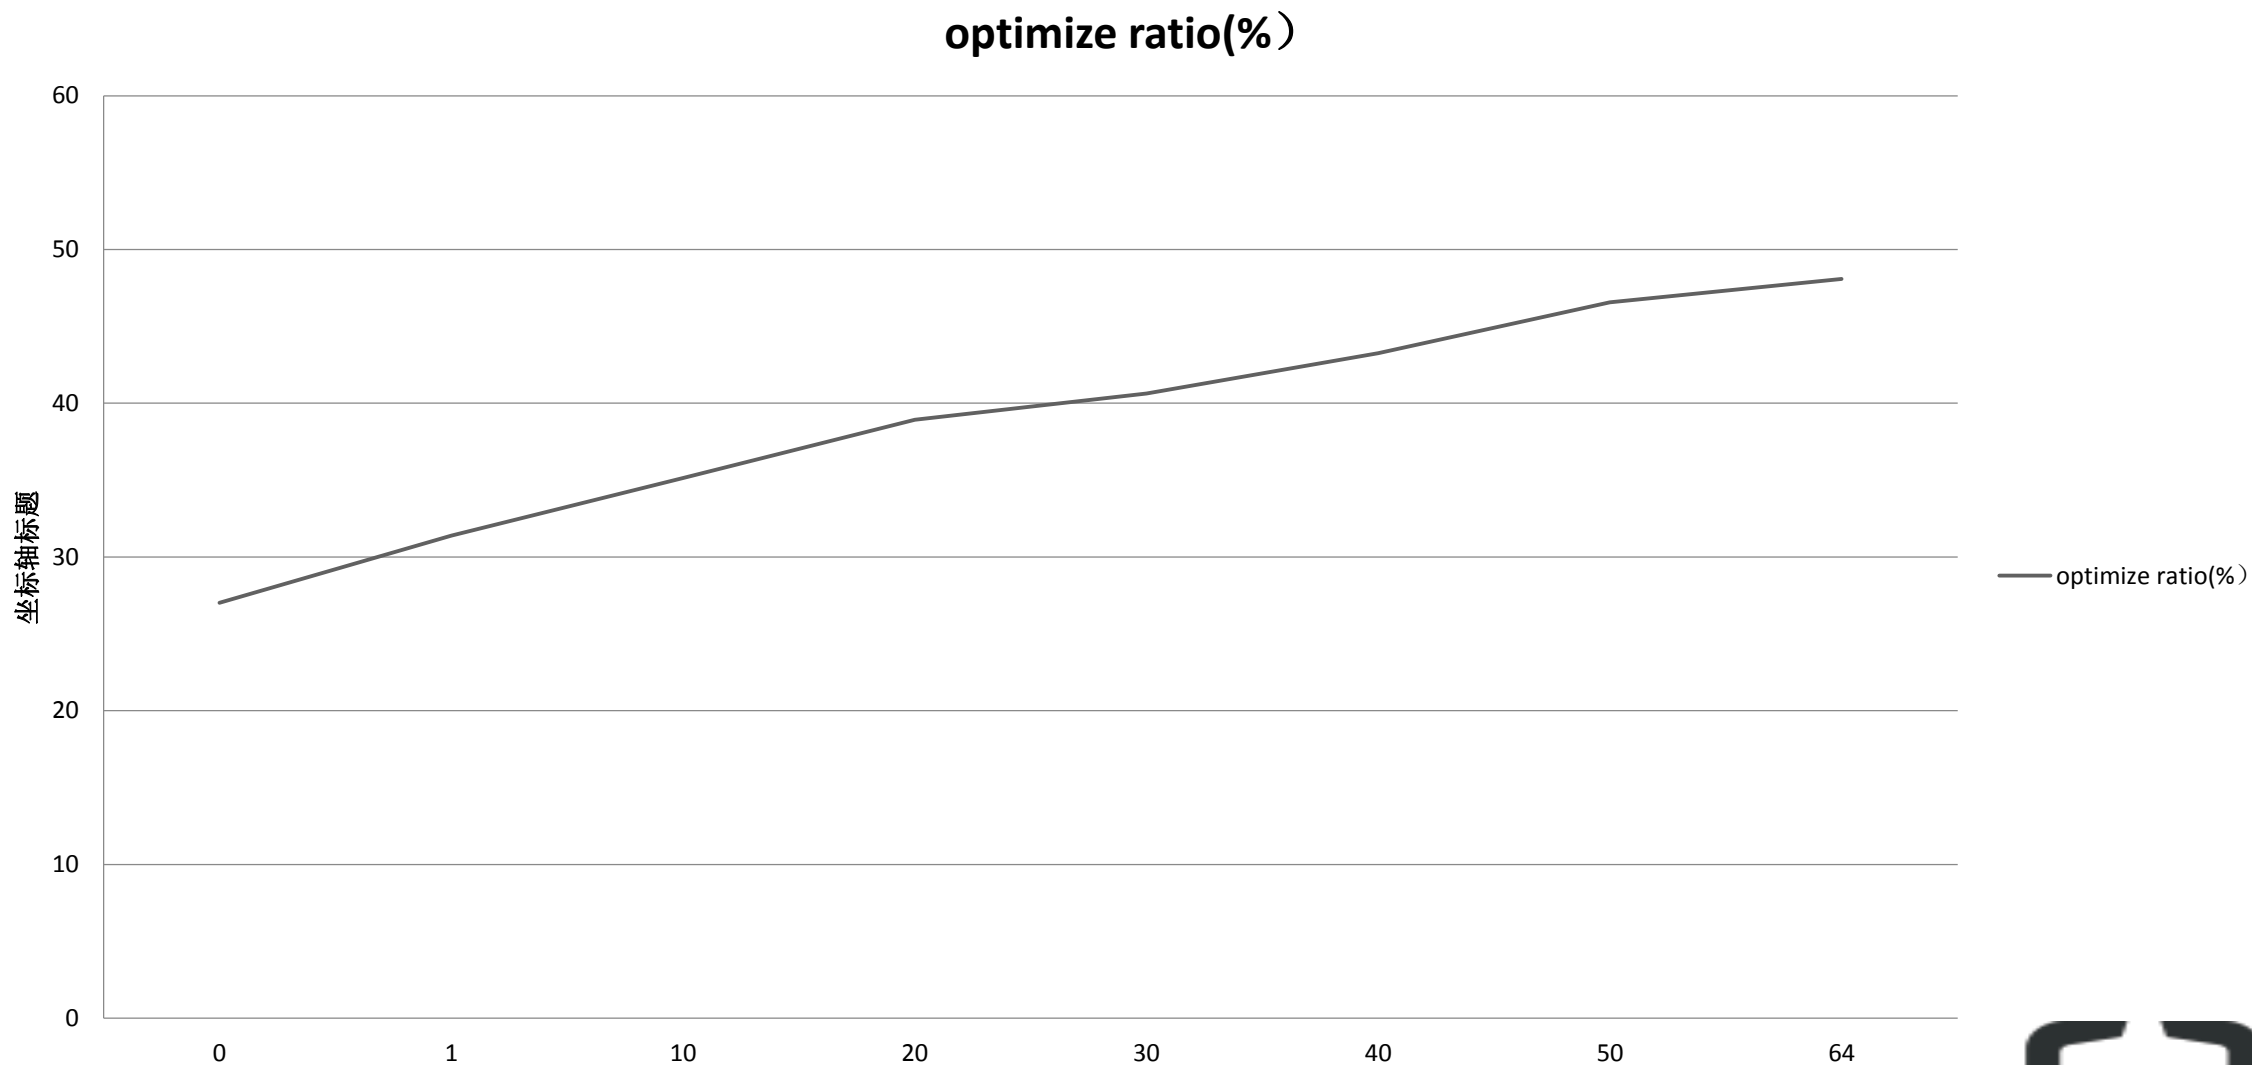
optimize time(单位 μ s)



```
SELECT * FROM t  
WHERE col1 > 1  
and ...  
and colN > 1
```




简单查询优化时间占执行时间的比例分析




optimize time(单位 μ s)

```
SELECT * FROM t1 .. tN  
WHERE t1.col1 = tN.col1  
and ...  
and t1.colN = tN.colN
```

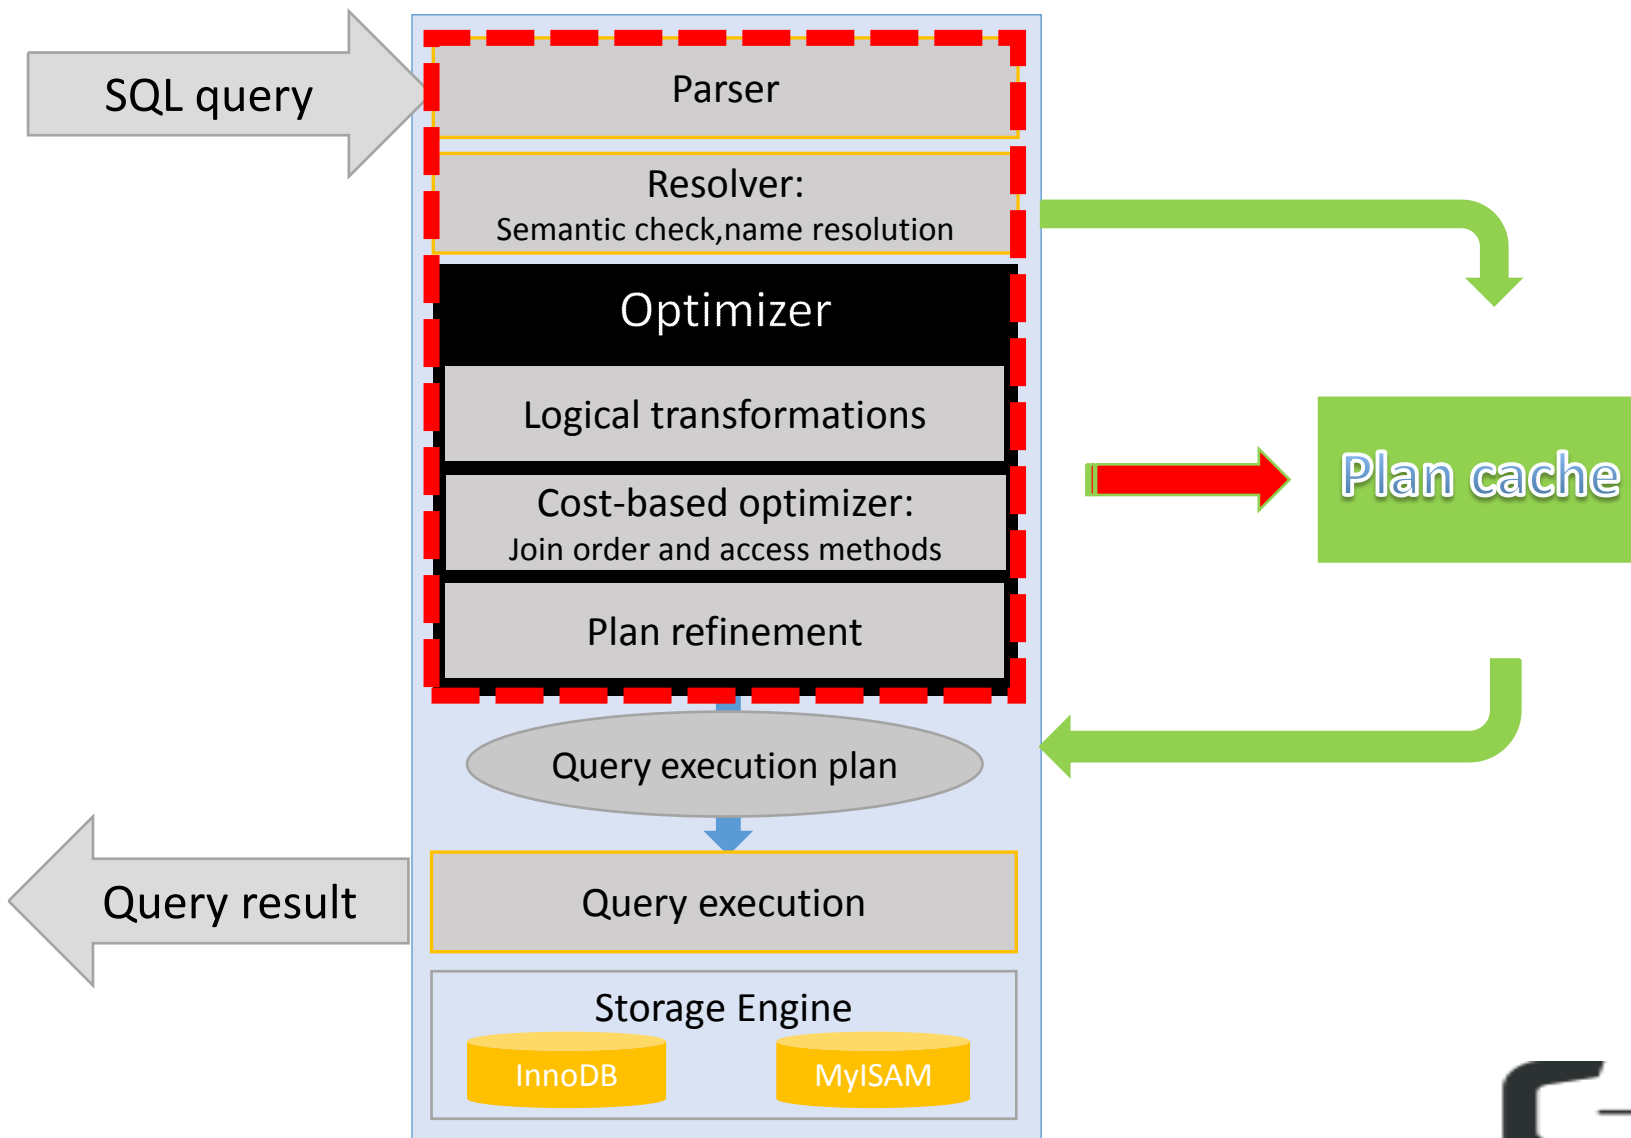


多表连接优化时间占执行时间的比例分析

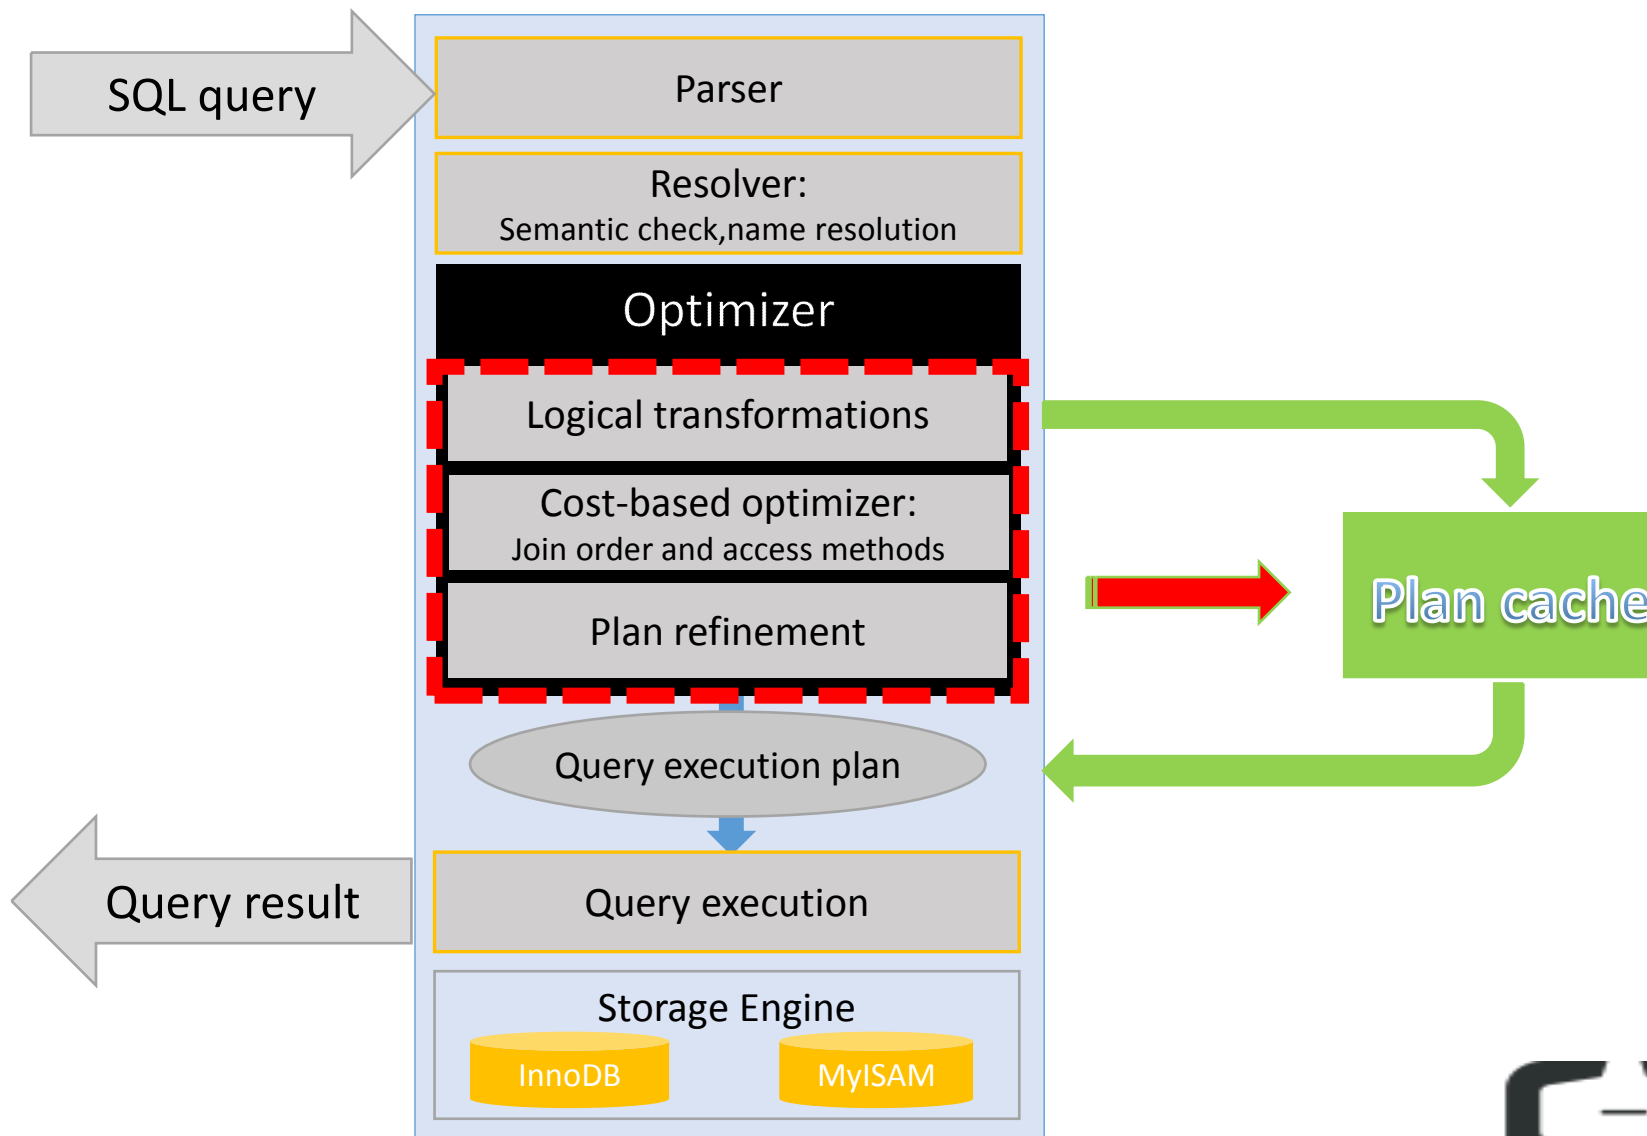
optimize ratio(%)



AliSQL Plan cache的实现



AliSQL Plan cache的实现



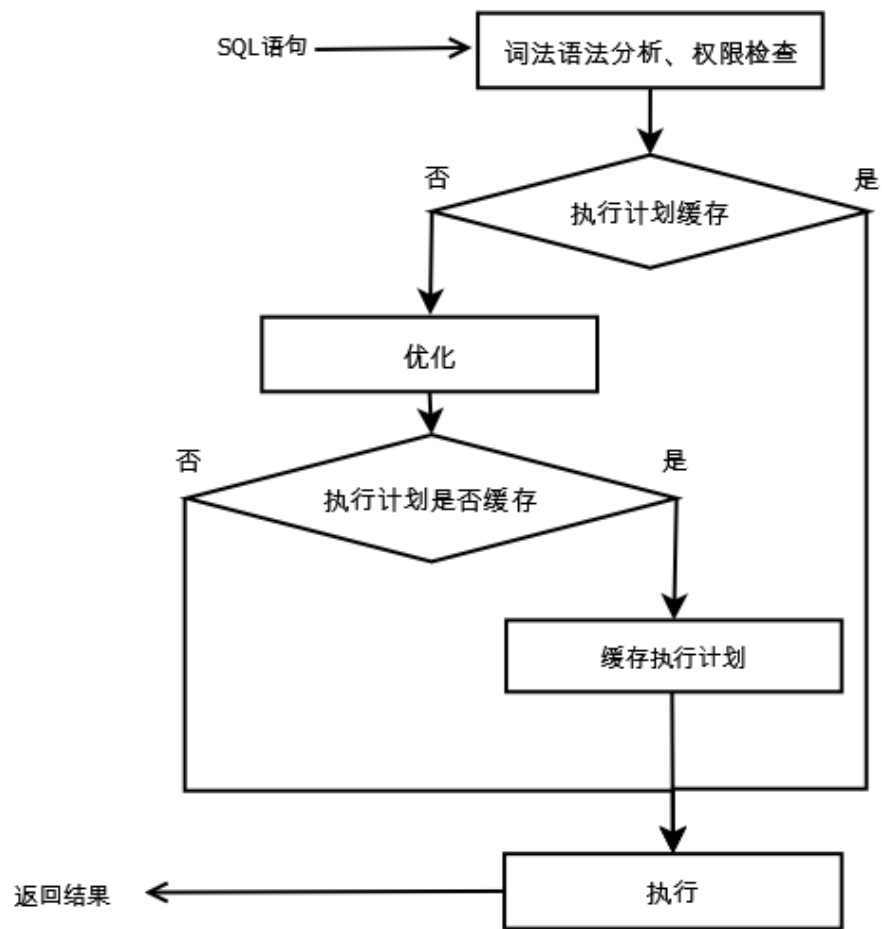


图1

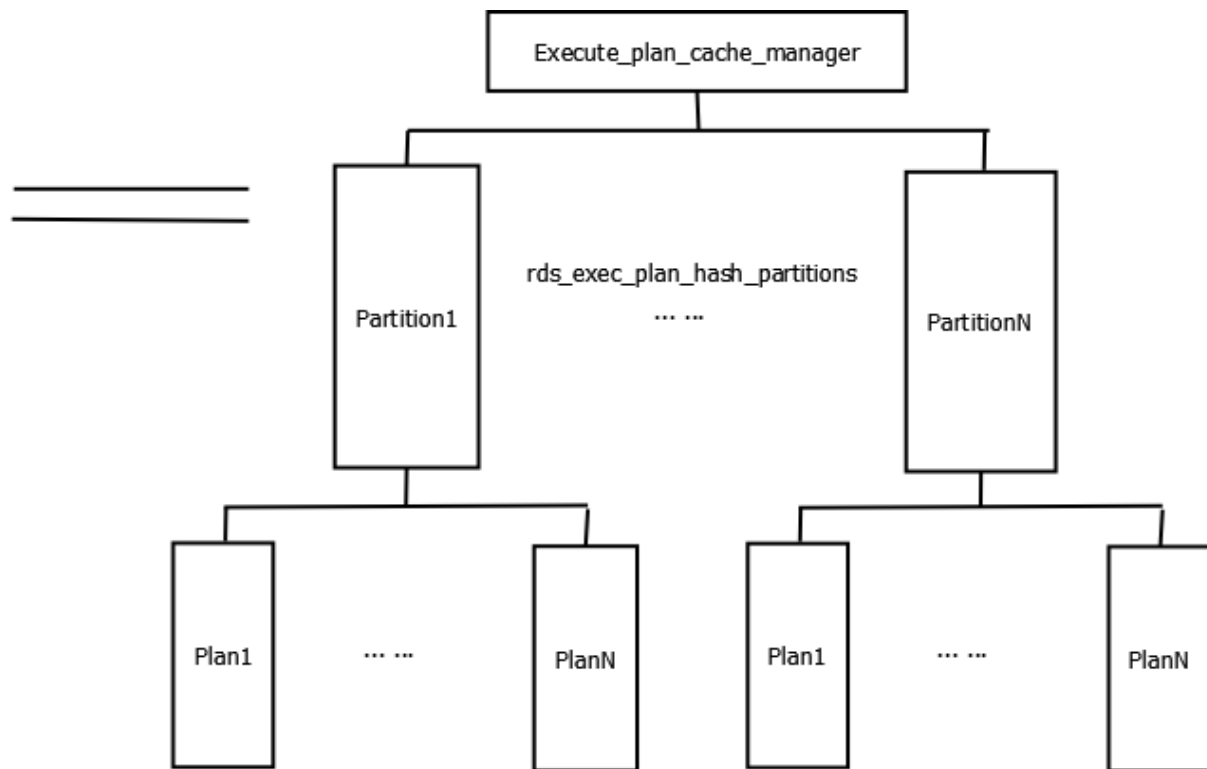


图2



变量名	注释
rds_enable_sql_digest	a) ON 打开sql_digest功能。b) OFF 关闭sql_digest功能(默认)
rds_max_digest_length	a) 设置SQL语句中常量替换后的的长度，设置范围是（128 ~ 1M）。b) 默认值是4K。c) 如果SQL语句长度大于该值，query的plan不会被缓存。
rds_enable_exec_plan_cache	a) ON 打开Plan Cache功能。b) OFF 关闭Plan Cache功能，清空Plan Cache(默认)。
rds_exec_plan_hash_partitions	可以有多少个partition来并发存储query plan。默认值是8
rds_max_plan_cache_mem_size	如果Plan Cache中分配到某个Partition中的记录所使用的内存超过了rds_max_plan_cache_mem_size的平均数，即 $rds_max_plan_cache_mem_size / rds_exec_plan_hash_parititions$ ，Plan Cache将利用LRU对存在的执行计划记录进行淘汰。



1. TABLE change
2. Set global `rds_enable_exec_plan_cache = off`
3. Hint `NO_PLAN_CACHE`
4. Hint `FORCE_UPDATE_PLAN_CACHE`



1. information_schema.exec_cache_status

| SQL_PRINT | SQL_DIGEST | TABLE_NAME | KEYS | USED_MEMORY | HIT_COUNT | EXTENDED |

2. optimizer_trace

```
{
  "plan_cache": [
    {
      "table": "`tt`",
      "rows": 0,
      "cost": 0,
      "use_cached_plan": "yes",
      "scan type": "ALL"
    }
  ]
}
```

3. show status like ‘% plan_cache %’

```
mysql> show status like '%plan_cache%';
```

```
+-----+-----+
| Variable_name      | Value |
+-----+-----+
| Execute_plan_cache_hits      | 1     |
| Execute_plan_cache_records   | 1     |
| Execute_plan_cache_used_memory | 1168  |
+-----+-----+
```

3 rows in set (0.01 sec)

4. Debug log (--debug)



```
mysql> show create table t1;
+-----+
-----+
| Table | Create Table
+-----+
-----+
| t1    | CREATE TABLE `t1` (
  `a` int(11) DEFAULT NULL,
  `b` int(11) DEFAULT NULL,
  KEY `id1` (`a`),
  KEY `id2` (`b`),
  KEY `id3` (`a`,`b`)
) ENGINE=InnoDB DEFAULT CHARSET=utf8 |
+-----+
-----+
1 row in set (0.00 sec)

mysql> explain select * from t1 where a > 1;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE      | t1    | range | id1,id3       | id1 | 5        | NULL | 2    | Using index condition
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```



```
mysql> set global rds_enable_exec_plan_cache=on;  
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> select * from t1 where a > 1;
```

```
+-----+-----+  
| a     | b     |  
+-----+-----+  
| 2     | 3     |  
| 3     | 4     |  
+-----+-----+
```

```
2 rows in set (0.00 sec)
```

```
mysql> select * from information_schema.exec_cache_status where table_name like '%t1%';
```

```
+-----+-----+-----+-----+-----+-----+-----+  
| SQL_PRINT | SQL_DIGEST | TABLE_NAME | KEYS | USED_MEMORY | HIT_COUNT | EXTENDED |  
+-----+-----+-----+-----+-----+-----+-----+  
| 3122753265 | SELECT * FROM t1 WHERE a > ? | test.t1 | id1 | 160 | 0 | range, partition number: 1 |  
+-----+-----+-----+-----+-----+-----+-----+
```

```
1 row in set (0.00 sec)
```



AliSQL Plan cache的使用 – 示例

```
mysql> select * from t1 where a > 1;
```

```
+-----+-----+  
| a     | b     |  
+-----+-----+  
|      2 |      3 |  
|      3 |      4 |  
+-----+-----+
```

```
2 rows in set (0.00 sec)
```

```
mysql> select * from information_schema.exec_cache_status where table_name like '%t1%';
```

```
+-----+-----+-----+-----+-----+-----+-----+  
| SQL_PRINT | SQL_DIGEST | TABLE_NAME | KEYS | USED_MEMORY | HIT_COUNT | EXTENDED |  
+-----+-----+-----+-----+-----+-----+-----+  
| 3122753265 | SELECT * FROM t1 WHERE a > ? | test.t1 | id1 | 160 | 1 | range, partition number: 1 |  
+-----+-----+-----+-----+-----+-----+-----+
```

```
1 row in set (0.00 sec)
```

```
mysql> show status like '%plan_cache_hits%';
```

```
+-----+-----+  
| Variable_name | Value |  
+-----+-----+  
| Execute_plan_cache_hits | 1 |  
+-----+-----+
```

```
1 row in set (0.00 sec)
```



AliSQL Plan cache的使用 – 示例

```
mysql> SELECT NO_PLAN_CACHE * FROM t1 where a > 1;
+-----+-----+
| a     | b     |
+-----+-----+
| 2     | 3     |
| 3     | 4     |
+-----+-----+
2 rows in set (0.01 sec)

mysql> SELECT * FROM information_schema.exec_cache_status;
+-----+-----+-----+-----+-----+-----+-----+
| SQL_PRINT | SQL_DIGEST | TABLE_NAME | KEYS | USED_MEMORY | HIT_COUNT | EXTENDED |
+-----+-----+-----+-----+-----+-----+-----+
| 3417261669 | SELECT * FROM `t1` WHERE `a` > ? | test.t1 | id1 | 160 | 1 | range, partition number: 5 |
+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)

mysql> DELETE FROM t1 WHERE a = 1;
Query OK, 10 rows affected (0.03 sec)


mysql> EXPLAIN SELECT * FROM t1 where a > 1;
+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | t1 | index | id1,id3 | id3 | 10 | NULL | 1 | Using where; Using index |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)

mysql> SELECT FORCE_UPDATE_PLAN_CACHE * FROM t1 where a > 1;
+-----+-----+
| a     | b     |
+-----+-----+
| 2     | 3     |
| 3     | 4     |
+-----+-----+
2 rows in set (0.00 sec)

mysql> SELECT * FROM information_schema.exec_cache_status;
+-----+-----+-----+-----+-----+-----+-----+
| SQL_PRINT | SQL_DIGEST | TABLE_NAME | KEYS | USED_MEMORY | HIT_COUNT | EXTENDED |
+-----+-----+-----+-----+-----+-----+-----+
| 3417261669 | SELECT * FROM `t1` WHERE `a` > ? | test.t1 | id3 | 160 | 0 | index, partition number: 5 |
+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```




Unit: QPS




QPS improvement



Unit: second



Read-only



Read-write



- RDS_5616
- PolarDB
- RDS 金融版
- AliSQL即将开源，欢迎大家试用

注： 目前控制台不支持Plan Cache设置，需要联系DBA



- 当前已经支持的场景

- SELECT 查询
- 单表

- 即将支持的场景：

- 支持多表连接
- 自动化管理Plan cache
 - Statistics
 - Optimizer_switch
- 支持 Plan Cache持久化
- 通过Plan Cache直接调整执行计划



Q & A



THANK YOU!

