

Details for Query 2

Submitted Time: 2023/11/07 22:53:04

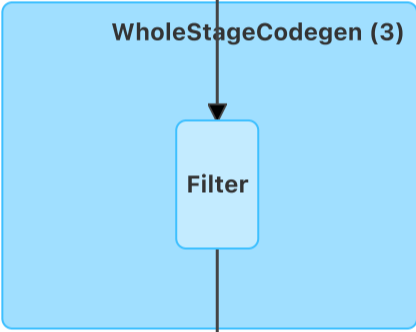
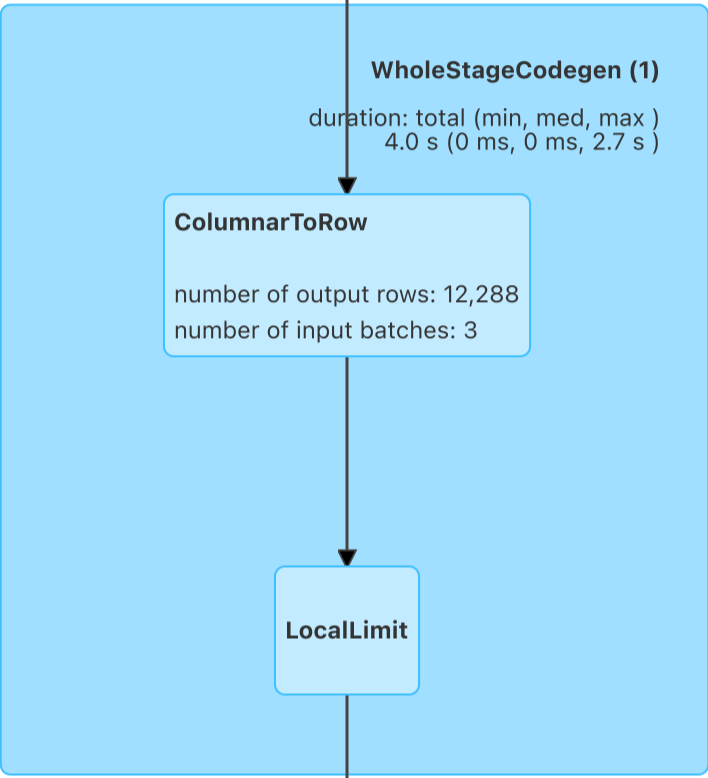
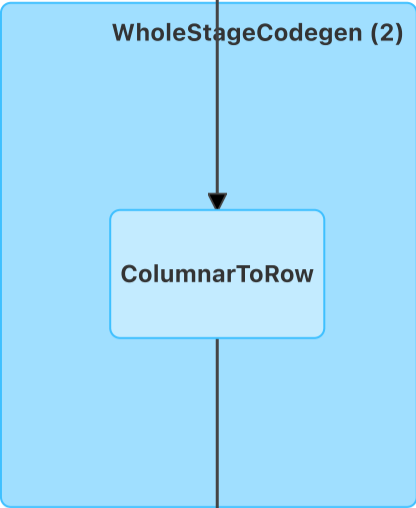
Duration: 11 s

Running Jobs: [0 1](#)

Show the Stage ID and Task ID that corresponds to the max metric

Scan orc fin_dm_data_ai.dm_ads_prea_all_score_fusion_6_7
number of files read: 2,048
metadata time: 5 ms
size of files read: 181.9 GiB

Scan orc fin_dm_data_ai.dm_ads_fkblack_wj_all
number of files read: 166
scan time total (min, med, max)
3.7 s (0 ms, 0 ms, 2.4 s)
metadata time: 5 ms
size of files read: 10.9 GiB
number of output rows: 12,288



Exchange
shuffle records written: 63
shuffle write time total (min, med, max)
0 ms (0 ms, 0 ms, 0 ms)
data size total (min, med, max)
3.4 KiB (0.0 B, 0.0 B, 1176.0 B)
number of partitions: 500
shuffle bytes written total (min, med, max)
2.4 KiB (0.0 B, 0.0 B, 842.0 B)

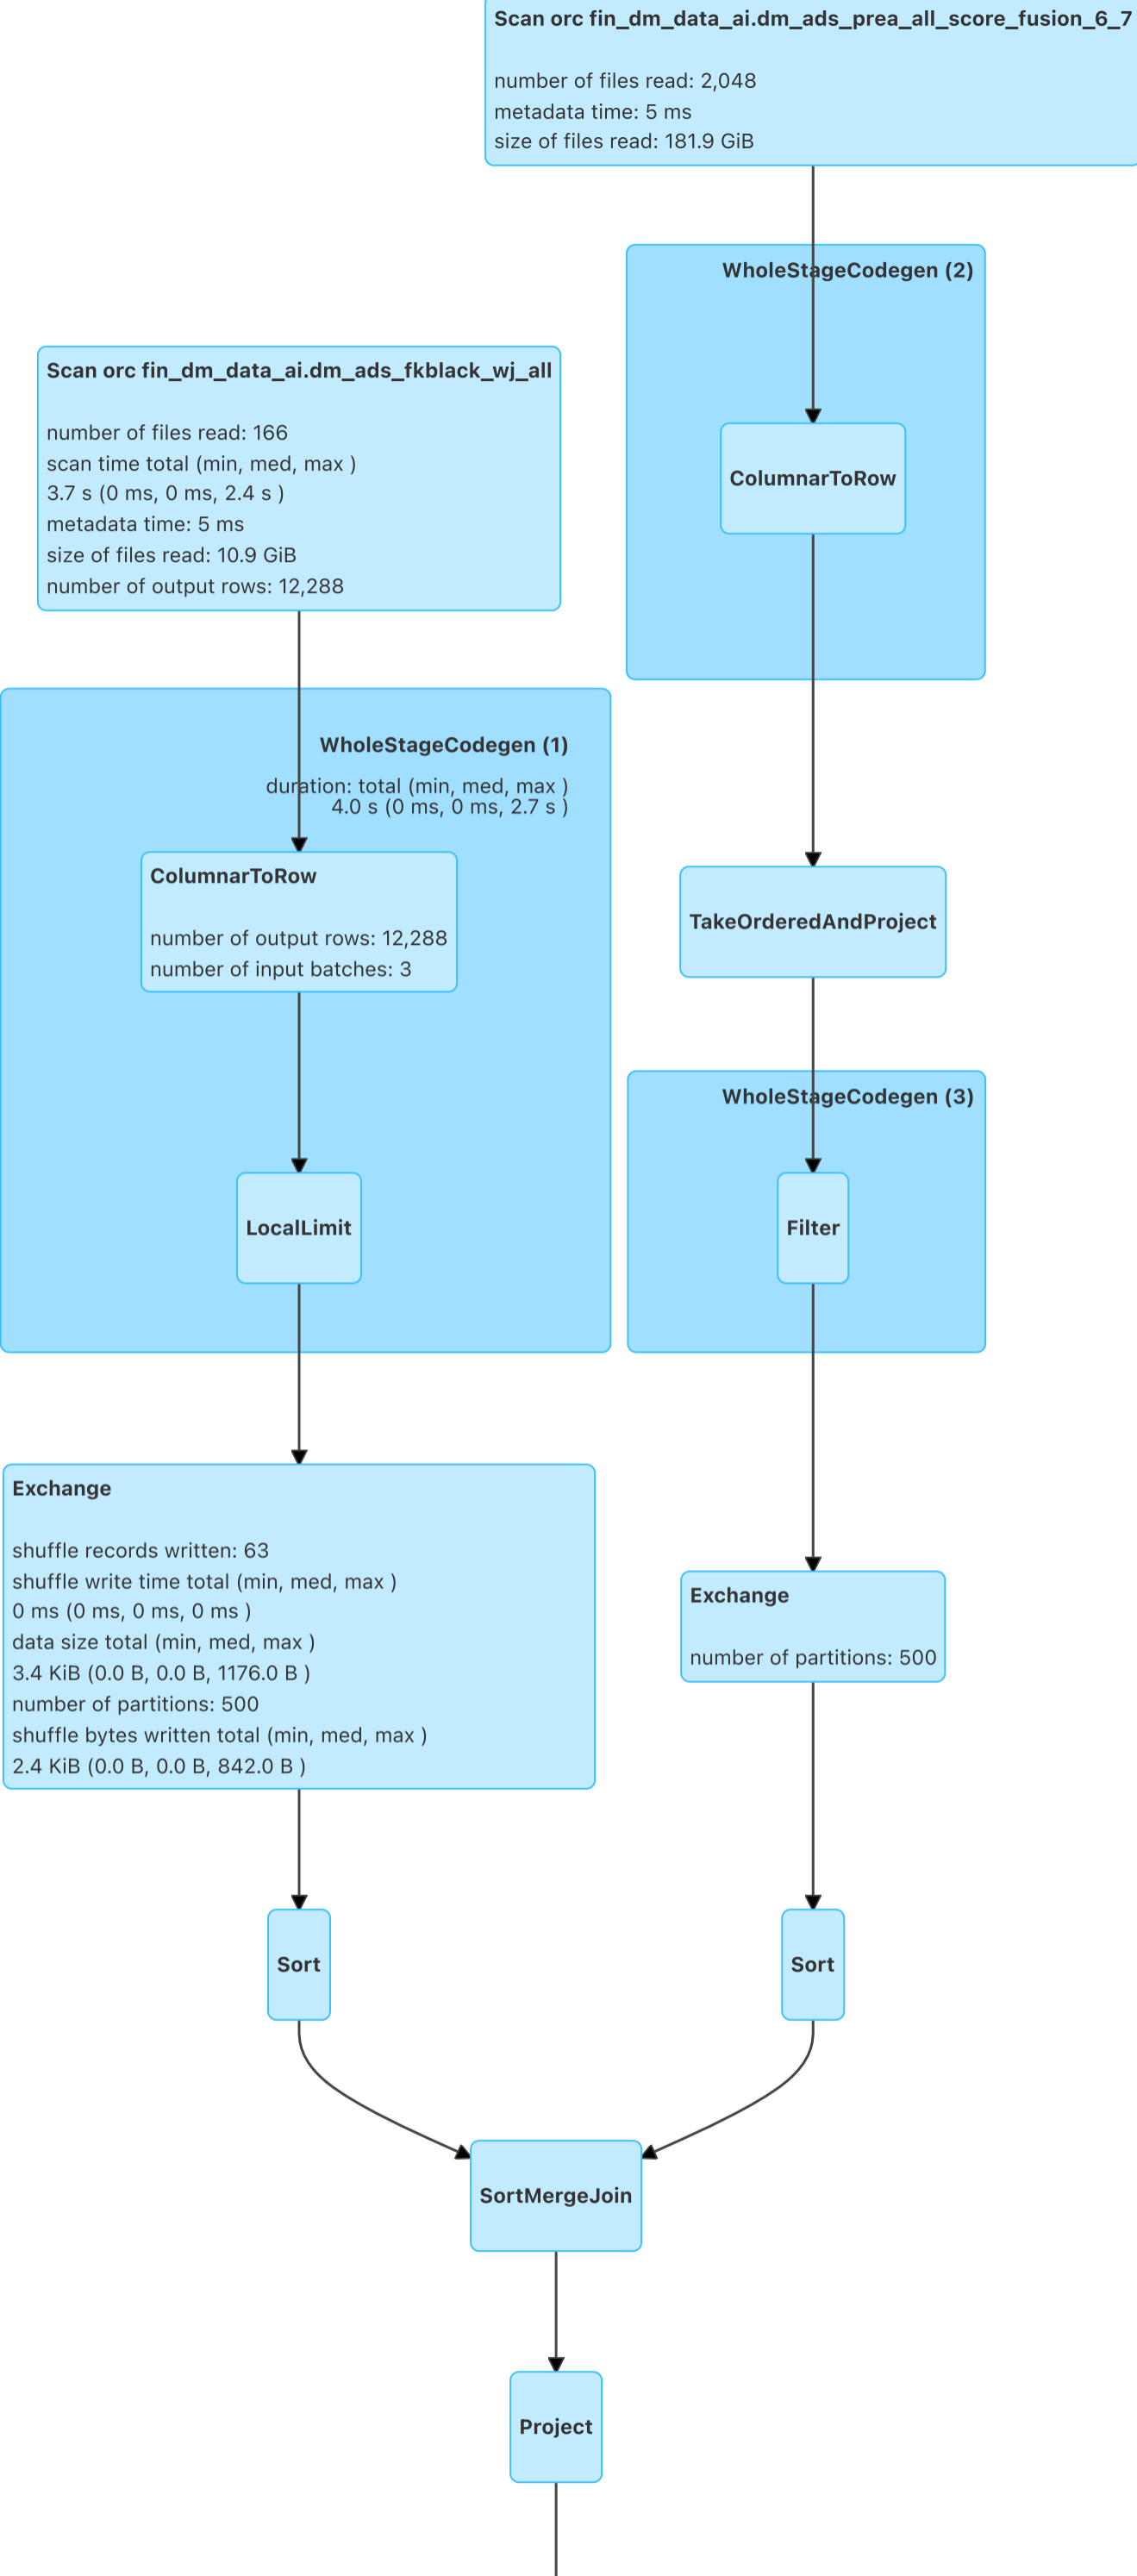
Exchange
number of partitions: 500


Sort

Sort

SortMergeJoin

Project





▼ Details

```

== Physical Plan ==
AdaptiveSparkPlan (27)
+- == Current Plan ==
  CollectLimit (16)
  +- Project (15)
    +- SortMergeJoin LeftOuter (14)
      :- Sort (6)
      : +- ShuffleQueryStage (5)
      :   +- Exchange (4)
      :     +- * LocalLimit (3)
      :     +- * ColumnarToRow (2)
      :     +- Scan orc fin_dm_data_ai.dm_ads_fkblack_wj_all (1)
    +- Sort (13)
      +- ShuffleQueryStage (12)
        +- Exchange (11)
          +- * Filter (10)
            +- TakeOrderedAndProject (9)
              +- * ColumnarToRow (8)
                +- Scan orc fin_dm_data_ai.dm_ads_prea_all_score_fusion_6_7 (7)

```

```

+- == Initial Plan ==
  CollectLimit (26)
  +- Project (25)
    +- SortMergeJoin LeftOuter (24)
      :- Sort (19)
      : +- Exchange (18)
      :   +- LocalLimit (17)
      :   +- Scan orc fin_dm_data_ai.dm_ads_fkblack_wj_all (1)
    +- Sort (23)
      +- Exchange (22)
      +- Filter (21)
        +- TakeOrderedAndProject (20)
          +- Scan orc fin_dm_data_ai.dm_ads_prea_all_score_fusion_6_7 (7)

```

(1) Scan orc fin_dm_data_ai.dm_ads_fkblack_wj_all

Output [2]: [deviceid#6, score#7]

Batched: true

Location: InMemoryFileIndex [hdfs://360jinronglycc/user/hive/warehouse/fin_dm_data_ai.db/dm_ads_fkblack_wj_all]

ReadSchema: struct<deviceid:string,score:double>

(2) ColumnarToRow [codegen id : 1]

Input [2]: [deviceid#6, score#7]

(3) LocalLimit [codegen id : 1]

Input [2]: [deviceid#6, score#7]

Arguments: 21

(4) Exchange

Input [2]: [deviceid#6, score#7]

Arguments: hashpartitioning(deviceid#6, 500), ENSURE_REQUIREMENTS, [plan_id=56]

(5) ShuffleQueryStage

Output [2]: [deviceid#6, score#7]

Arguments: 0

(6) Sort

Input [2]: [deviceid#6, score#7]

Arguments: [deviceid#6 ASC NULLS FIRST], false, 0

(7) Scan orc fin_dm_data_ai.dm_ads_prea_all_score_fusion_6_7

Output [2]: [deviceid#0, score#1]

Batched: true

Location: InMemoryFileIndex [hdfs://360jinronglycc/user/hive/warehouse/fin_dm_data_ai.db/dm_ads_prea_all_score_fusion_6_7]

ReadSchema: struct<deviceid:string,score:float>

(8) ColumnarToRow [codegen id : 2]

Input [2]: [deviceid#0, score#1]

(9) TakeOrderedAndProject

Input [2]: [deviceid#0, score#1]

Arguments: 1000000, [score#1 ASC NULLS FIRST], [deviceid#0, score#1]

(10) Filter [codegen id : 3]

Input [2]: [deviceid#0, score#1]

Condition : isnotnull(deviceid#0)

(11) Exchange

Input [2]: [deviceid#0, score#1]

Arguments: hashpartitioning(deviceid#0, 500), ENSURE_REQUIREMENTS, [plan_id=83]

(12) ShuffleQueryStage

Output [2]: [deviceid#0, score#1]

Arguments: 1

```
(13) Sort
Input [2]: [deviceid#0, score#1]
Arguments: [deviceid#0 ASC NULLS FIRST], false, 0

(14) SortMergeJoin
Left keys [1]: [deviceid#6]
Right keys [1]: [deviceid#0]
Join condition: None

(15) Project
Output [4]: [deviceid#6, cast(score#7 as string) AS score#21, deviceid#0, cast(score#1 as string) AS score#23]
Input [4]: [deviceid#6, score#7, deviceid#0, score#1]

(16) CollectLimit
Input [4]: [deviceid#6, score#21, deviceid#0, score#23]
Arguments: 21

(17) LocalLimit
Input [2]: [deviceid#6, score#7]
Arguments: 21

(18) Exchange
Input [2]: [deviceid#6, score#7]
Arguments: hashpartitioning(deviceid#6, 500), ENSURE_REQUIREMENTS, [plan_id=38]

(19) Sort
Input [2]: [deviceid#6, score#7]
Arguments: [deviceid#6 ASC NULLS FIRST], false, 0

(20) TakeOrderedAndProject
Input [2]: [deviceid#0, score#1]
Arguments: 1000000, [score#1 ASC NULLS FIRST], [deviceid#0, score#1]

(21) Filter
Input [2]: [deviceid#0, score#1]
Condition : isnotnull(deviceid#0)

(22) Exchange
Input [2]: [deviceid#0, score#1]
Arguments: hashpartitioning(deviceid#0, 500), ENSURE_REQUIREMENTS, [plan_id=39]

(23) Sort
Input [2]: [deviceid#0, score#1]
Arguments: [deviceid#0 ASC NULLS FIRST], false, 0

(24) SortMergeJoin
Left keys [1]: [deviceid#6]
Right keys [1]: [deviceid#0]
Join condition: None

(25) Project
Output [4]: [deviceid#6, cast(score#7 as string) AS score#21, deviceid#0, cast(score#1 as string) AS score#23]
Input [4]: [deviceid#6, score#7, deviceid#0, score#1]

(26) CollectLimit
Input [4]: [deviceid#6, score#21, deviceid#0, score#23]
Arguments: 21

(27) AdaptiveSparkPlan
Output [4]: [deviceid#6, score#21, deviceid#0, score#23]
Arguments: isFinalPlan=false
```