

Metrics_SparkUI_projet

Voici toutes les captures d'écran des métriques avec le tableau final

- **Q1 baseline**

↓

Exchange

shuffle records written: 6,072,131
remote merged reqs duration: 0 ms
remote merged blocks fetched: 0
records read: 6,072,131
local bytes read total (min, med, max)
278.7 MiB (11.1 MiB, 16.7 MiB, 16.8 MiB)
merged fetch fallback count: 0
local blocks read: 272
remote merged chunks fetched: 0
remote blocks read: 0
data size total (min, med, max)
512.3 MiB (24.7 MiB, 32.5 MiB, 32.6 MiB)
local merged bytes read: 0.0 B
local merged chunks fetched: 0
shuffle write time total (min, med, max)
6.2 s (184 ms, 403 ms, 490 ms)
remote merged bytes read: 0.0 B
local merged blocks fetched: 0
corrupt merged block chunks: 0
fetch wait time total (min, med, max)
0 ms (0 ms, 0 ms, 0 ms)
remote bytes read: 0.0 B
number of partitions: 200
remote reqs duration: 0 ms
remote bytes read to disk: 0.0 B
shuffle bytes written total (min, med, max)
278.7 MiB (13.5 MiB, 17.7 MiB, 17.7 MiB)

↓

Scan csv

number of output rows: 6,072,136
number of files read: 1
metadata time: 0 ms
size of files read: 261.1 MiB

Submitted Time: 2026/01/04 01:12:26

Duration: 12 s

Succeeded Jobs: 61 62 63

- **Q2 baseline**

↓

Exchange

shuffle records written: 6,072,131
remote merged reqs duration: 0 ms
remote merged blocks fetched: 0
records read: 6,072,131
local bytes read total (min, med, max)
278.7 MiB (11.1 MiB, 16.7 MiB, 16.8 MiB)
merged fetch fallback count: 0
local blocks read: 272
remote merged chunks fetched: 0
remote blocks read: 0
data size total (min, med, max)
512.3 MiB (24.7 MiB, 32.5 MiB, 32.6 MiB)
local merged bytes read: 0.0 B
local merged chunks fetched: 0
shuffle write time total (min, med, max)
2.2 s (58 ms, 99 ms, 296 ms)
remote merged bytes read: 0.0 B
local merged blocks fetched: 0
corrupt merged block chunks: 0
fetch wait time total (min, med, max)
0 ms (0 ms, 0 ms, 0 ms)
remote bytes read: 0.0 B
number of partitions: 200
remote reqs duration: 0 ms
remote bytes read to disk: 0.0 B
shuffle bytes written total (min, med, max)
278.7 MiB (13.5 MiB, 17.7 MiB, 17.7 MiB)

↓

Scan csv

number of output rows: 6,072,136
number of files read: 1
metadata time: 0 ms
size of files read: 261.1 MiB

Submitted Time: 2026/01/04 14:18:44

Duration: 16 s

Succeeded Jobs: 28 29 30

- **Q3 baseline**

Exchange

shuffle records written: 6,072,131
remote merged reqs duration: 0 ms
remote merged blocks fetched: 0
records read: 6,072,131
local bytes read total (min, med, max)
278.7 MiB (11.1 MiB, 16.7 MiB, 16.8 MiB)
merged fetch fallback count: 0
local blocks read: 272
remote merged chunks fetched: 0
remote blocks read: 0
data size total (min, med, max)
512.3 MiB (24.7 MiB, 32.5 MiB, 32.6 MiB)
local merged bytes read: 0.0 B
local merged chunks fetched: 0
shuffle write time total (min, med, max)
7.5 s (216 ms, 475 ms, 578 ms)
remote merged bytes read: 0.0 B
local merged blocks fetched: 0
corrupt merged block chunks: 0
fetch wait time total (min, med, max)
0 ms (0 ms, 0 ms, 0 ms)
remote bytes read: 0.0 B
number of partitions: 200
remote reqs duration: 0 ms
remote bytes read to disk: 0.0 B
shuffle bytes written total (min, med, max)
278.7 MiB (13.5 MiB, 17.7 MiB, 17.7 MiB)

Scan csv

number of output rows: 6,072,136
number of files read: 1
metadata time: 0 ms
size of files read: 261.1 MiB

Submitted Time: 2026/01/04 14:26:27

Duration: 11 s

Succeeded Jobs: 31 32 33

- **Q1 Optimized**



Exchange

shuffle records written: 6,072,131
remote merged reqs duration: 0 ms
remote merged blocks fetched: 0
records read: 6,072,131
local bytes read total (min, med, max)
122.3 MiB (622.4 KiB, 626.0 KiB, 630.6 KiB)
merged fetch fallback count: 0
local blocks read: 3,400
remote merged chunks fetched: 0
remote blocks read: 0
data size total (min, med, max)
251.3 MiB (10.0 MiB, 15.1 MiB, 15.1 MiB)
local merged bytes read: 0.0 B
local merged chunks fetched: 0
shuffle write time total (min, med, max)
3.9 s (16 ms, 243 ms, 323 ms)
remote merged bytes read: 0.0 B
local merged blocks fetched: 0
corrupt merged block chunks: 0
fetch wait time total (min, med, max)
0 ms (0 ms, 0 ms, 0 ms)
remote bytes read: 0.0 B
number of partitions: 200
remote reqs duration: 0 ms
remote bytes read to disk: 0.0 B
shuffle bytes written total (min, med, max)
122.3 MiB (5.0 MiB, 7.3 MiB, 7.4 MiB)



Scan parquet

number of files read: 17
scan time total (min, med, max)
11.5 s (277 ms, 702 ms, 777 ms)
metadata time: 0 ms
size of files read: 132.9 MiB
number of output rows: 6,072,131



Submitted Time: 2026/01/04 15:11:20

Duration: 7 s

Succeeded Jobs: 37 38 39

- **Q2 Optimized**

[plan_id=1978]

Exchange

shuffle records written: 6,072,131
remote merged reqs duration: 0 ms
remote merged blocks fetched: 0
records read: 6,072,131
local bytes read total (min, med, max)
271.7 MiB (1385.2 KiB, 1391.3 KiB, 1397.7 KiB)
merged fetch fallback count: 0
local blocks read: 3,400
remote merged chunks fetched: 0
remote blocks read: 0
data size total (min, med, max)
419.6 MiB (16.7 MiB, 25.2 MiB, 25.3 MiB)
local merged bytes read: 0.0 B
local merged chunks fetched: 0
shuffle write time total (min, med, max)
8.9 s (26 ms, 571 ms, 637 ms)
remote merged bytes read: 0.0 B
local merged blocks fetched: 0
corrupt merged block chunks: 0
fetch wait time total (min, med, max)
0 ms (0 ms, 0 ms, 0 ms)
remote bytes read: 0.0 B
number of partitions: 200
remote reqs duration: 0 ms
remote bytes read to disk: 0.0 B
shuffle bytes written total (min, med, max)
271.7 MiB (10.8 MiB, 16.3 MiB, 16.4 MiB)

Details for Query 19

Submitted Time: 2026/01/04 15:11:26

Duration: 13 s

Succeeded Jobs: [40](#) [41](#) [42](#)

▼ [Plan Visualization](#)

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

Scan parquet

number of files read: 17
scan time total (min, med, max)
19.5 s (644 ms, 1.2 s, 1.4 s)
metadata time: 1 ms
size of files read: 132.9 MiB
number of output rows: 6,072,131

- Q3 Optimized

↓

Exchange

shuffle records written: 6,072,131
 remote merged reqs duration: 0 ms
 remote merged blocks fetched: 0
 records read: 6,072,131
 local bytes read total (min, med, max)
 16.6 MiB (84.0 KiB, 84.8 KiB, 85.6 KiB)
 merged fetch fallback count: 0
 local blocks read: 3,400
 remote merged chunks fetched: 0
 remote blocks read: 0
 data size total (min, med, max)
 185.3 MiB (7.4 MiB, 11.1 MiB, 11.2 MiB)
 local merged bytes read: 0.0 B
 local merged chunks fetched: 0
 shuffle write time total (min, med, max)
 306 ms (8 ms, 16 ms, 33 ms)
 remote merged bytes read: 0.0 B
 local merged blocks fetched: 0
 corrupt merged block chunks: 0
 fetch wait time total (min, med, max)
 0 ms (0 ms, 0 ms, 0 ms)
 remote bytes read: 0.0 B
 number of partitions: 200
 remote reqs duration: 0 ms
 remote bytes read to disk: 0.0 B
 shuffle bytes written total (min, med, max)
 16.6 MiB (784.0 KiB, 1011.9 KiB, 1020.5 KiB)

Details for Query 20

Submitted Time: 2026/01/04 15:11:40

Duration: 3 s

Succeeded Jobs: 43 44 45

▼ [Plan Visualization](#)

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

Scan parquet

number of files read: 17
 scan time total (min, med, max)
 7.7 s (209 ms, 461 ms, 495 ms)
 metadata time: 0 ms
 size of files read: 132.9 MiB
 number of output rows: 6,072,131

- Tableau final

A	B	C	D	E	F	G	H	I	J	K
run_id	query	phase	files_read	input_size_bytes	shuffle_read_bytes	shuffle_write_bytes	etapsed_ms	notes	timestamp	
r1	Q1	baseline	1	512.3 MiB	278.7 MiB	278.7 MiB	12000	6M shuffle records	04/01/2026 01:12	
r1	Q1	optimized	17	132.9 MiB	122.3 MiB	122.3 MiB	7000		04/01/2026 15:11	
r1	Q2	baseline	1	512.3 MiB	278.7 MiB	278.7 MiB	16000	6M shuffle records	04/01/2026 14:18	
r1	Q2	optimized	17	132.9 MiB	271.7 MiB	271.7 MiB	13000	optimized layout, repartition	04/01/2026 15:11	
r1	Q3	baseline	1	512.3 MiB	278.7 MiB	278.7 MiB	11000	6M shuffle records	04/01/2026 14:26	
r1	Q3	optimized	17	132.9 MiB	16.6 MiB	16.6 MiB	3000		04/01/2026 15:11	