

Skewed Joins



Objective

Identify non-uniform data distribution

Observe and fix performance impact



To Remember

Non-uniform data distribution = data skews

Can cause massive performance problems

- extremely long jobs
- sometimes executor OOMs

Reason

- prior to join, Spark shuffles data
- same key will stay on the same executor
- if one key is disproportionate, that executor will have a disproportionately large task
- that executor takes longer than everyone else
- the whole job waits for that task to finish

Solution: include extra information in the join key set

- new key to join with
- redistribution of data by n+1 join keys
- uniform tasks

Spark rocks

