

Support Multiple Regions as Input to Each Mapper in MR Jobs

cloud-hbase@dev.qiyi.com

yeweichen@qiyi.com

Motivation

- For a Hadoop cluster, a job with a large HBase table as input always consumes a large amount of computing resources.
- For example, if we want to scan a table with 1000 regions, we need a job with 1000 mappers.
- This patch is to support one mapper using multiple regions as input.

Configuration

- Add a new configuration — ***hbase.mapreduce.scan.regionspermapper***
- ***hbase.mapreduce.scan.regionspermapper*** controls how many regions used as input for one mapper.
- For example, if we have a HBase table with 300 regions, and we set ***hbase.mapreduce.scan.regionspermapper*** = 3. Then we run a job to scan the table, the job will use only $300/3=100$ mappers.
- In this way, we can control the number of mappers using the following formula:
Number of Mappers = (Total region numbers) / `hbase.mapreduce.scan.regionspermapper`

Usage

- For example, if we have a table with 100 regions, the following table shows the relations between the configuration and the number of mappers.

hbase.mapreduce.scan.regionspermapper	Number of Mappers
1	$\text{ceil} (100 / 1) = 100$
2	$\text{ceil} (100 / 2) = 50$
3	$\text{ceil} (100 / 3) = 34$
5	$\text{ceil} (100 / 5) = 20$
10	$\text{ceil} (100 / 10) = 10$
30	$\text{ceil} (100 / 30) = 4$
99	$\text{ceil} (100 / 99) = 2$
100	$\text{ceil} (100 / 100) = 1$
150	$\text{ceil} (100 / 150) = 1$

Notice

- [HBase-2302](#) is a feature to exclude some specific regions from the MR job.
- When we set ***hbase.mapreduce.scan.regionspermapper > 1***, the feature in HBASE-2302 would not be supported. Because the mapper deal with only one TableSplit object, the multiple regions in one mapper must be continuous.