



MR-less mob compactions

- Goals:
 - Removed deleted mob objects
 - Avoid write amplification due to rewriting files
 - Skip rewriting if a threshold is not met
 - Allows for consolidation of mob hfiles (merge 8x 128mb file-> 1GB file)
 - Do not incur significantly large new memory requirements
 - Do not add 4gb memstore to get 4GB compacted mob hfiles.

Compact MOB Deletes Approach




Compact MOB Deletes Approach




- Normal Major Compaction writes out newly compacted hfile
- Also writes a DelMob hfile which contains references to compacted away mob refs
- Archives old HFiles

Compact MOB Deletes Approach



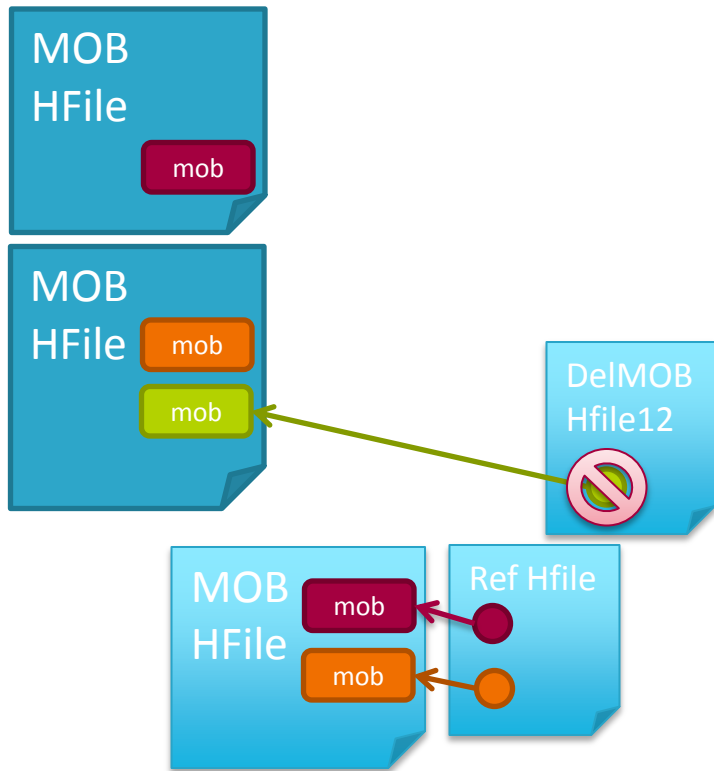
- Now we only need to consider mob compactions with del mobs files and mob files.
- No need to worry about major compactions

Compact MOB Deletes Approach



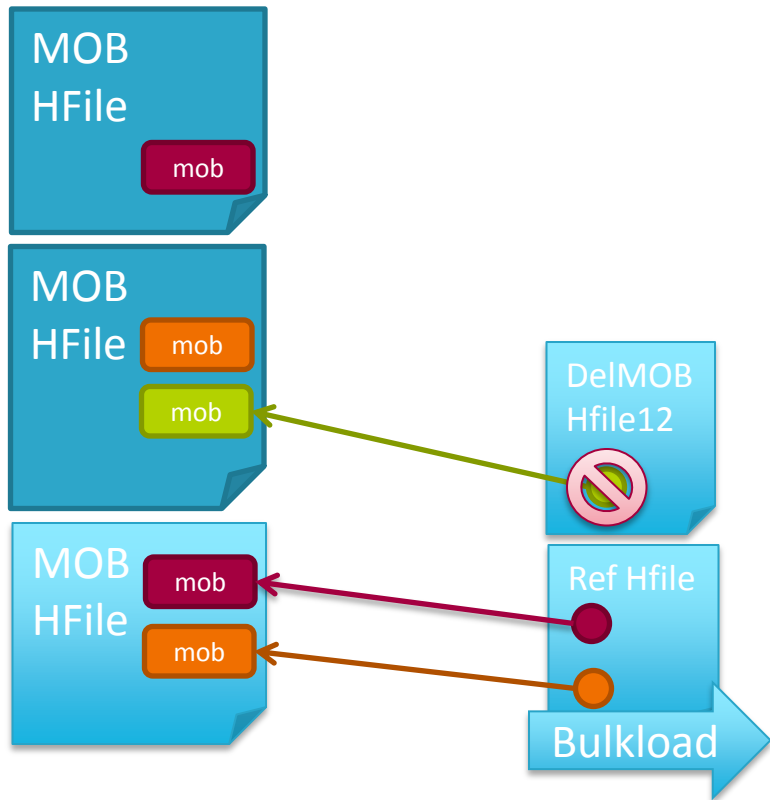
- Determine which MOB and DelMOB Hfiles to merge (scan all deletes or subrange of deletes)
- Select files where threshold is met via scan

Compact MOB Deletes Approach



- Compact selected mob files to tmp
- Creates new MOB Hfile and new ref hfile

Compact MOB Deletes Approach



- Move new MOB Hfile to mob dir.
- Bulk load ref hfile

Compact MOB Deletes Approach

- Archive original lobs

