Question
[reinhard lebensorger](#) · Sep 17, 2019

Cache and SQL Performance

We often have troubles with performance when joining tables with about 10 million rows together.

Sometimes we can't say what the cache engine is doing, it's just sucks and never finishes.

Then we have to rebuild the query and loop first on Orders and then on OrderPositions.

this is what we want:

EDITED: the join was only forgotten

select p.hpos,p.upos,p.spos, p.Fa,p.Auftnr,a.AA, p.art
from
auf.aufpos p ,Auf.Auftrag a
where a.Fa=1 and a.Fa=p.Fa
and a.AuftDat > ($HOROLOG -1000)
and a.Auftnr=p.Auftnr

this query does not finish...

then we rebuild and work with

select o.col1, o.col2,
from orders o
where o.orderDate > ($H -1000)

and then loop :o(

indizes are correct, all join cols are indexed

i can't believe that SQL isn't available for us when we join tables with "more" data because exactly for this purpose databases have been developed

maybe there are hints like on other databases like "UPDATE STATISTICS"

also: write your experience, do your queries also "fail" by never return?

my suspicion goes in direction queryoptimizer, maybe the system has wrong statistics or something else.

best regards, reinhard

#Caché