

Database - Feature #5453

optimize FindQuery hasOne and hasAny implementations

06/15/2021 12:09 PM - Eric Faulhaber

Status: New	Start date:
Priority: Normal	Due date:
Assignee:	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version:	version:
billable: No	
vendor_id: GCD	
Description	
Related issues:	
Related to Database - Feature #7076: make CAN-FIND avoid hydration Closed	

History

#1 - 06/15/2021 12:15 PM - Eric Faulhaber

We currently convert standalone CAN-FIND calls to use FindQuery.hasAny or FindQuery.hasOne. These methods will fetch the entire record and hydrate a DMO in order to answer the question of whether a record exists, then throw that DMO away. Under the right conditions, we should be able to skip the fetch and just perform a projection query (i.e., primary key only), which would be faster.

This will require some refactoring of RAQ.execute and RAQ.executeImpl. Complications to consider w.r.t. "the right conditions":

- the fast-find cache;
- the dirty database.

#2 - 06/15/2021 07:12 PM - Ovidiu Maxiniuc

Alternatively, we can keep the record in session cache, hoping a future query will have the result already fetched. Usually, the hasXXX() queries are used before issuing a bigger/more resource consuming query. If the query is not issued, the record is dropped, as per cache policy.

#3 - 01/26/2023 04:09 AM - Alexandru Lungu

- Related to Feature #7076: make CAN-FIND avoid hydration added