

Database - Bug #7993

Create the fast-copy mapping tables (MASTER_PK_MAPPING and BEFORE_PK_MAPPING) only once per session

11/01/2023 11:01 AM - Alexandru Lungu

Status:	Test	Start date:	
Priority:	Normal	Due date:	
Assignee:	Alexandru Lungu	% Done:	100%
Category:		Estimated time:	0.00 hour
Target version:		case_num:	
billable:	No		
vendor_id:	GCD		
Description			

History

#1 - 11/01/2023 11:38 AM - Alexandru Lungu

- Status changed from WIP to Review

- % Done changed from 0 to 100

The fast-copy process uses two intermediary tables (BEFORE__PK__MAPPING and MASTER__PK__MAPPING) to fulfill more complex copy operations including normalized extent tables or before tables. The tables store the PK mapping between the records in the master table and are reused by the extent/before copy attempts to properly remap to the new PK.

In 7993a, I made the required changes to have only **one** pair of BEFORE__PK__MAPPING and MASTER__PK__MAPPING per session **and** not per fast copy operation. In short, I replaced local CREATE TABLE and DROP TABLE with a global CREATE TABLE and local INSERT INTO + DELETE FROM (which can be transformed into a TRUNCATE at the right time - TODO).

The problem we had for a while was that by DROPPING the PK mapping tables, the statements were reforced to be reparsed (even if prepared). Now, because there is one single MASTER__PK__MAPPING and one single BEFORE__PK__MAPPING, fast-copy statements are no longer reparsed to honor schema changes.

I tested with a POC application - there was no regression. With SQL-logging, I've seen 50% less insert into select from statements in my >0.1ms execution time list. Please review.

I will attempt now to test with a bigger customer application and eventually profile.

#2 - 11/01/2023 11:07 PM - Ovidiu Maxiniuc

- Status changed from Review to Test

Review of c14813/7993a.

I am OK with the code. As described above, it should offer a bit of performance boost if the operation is performed multiple times. In singular cases there should be no difference.

I do not know the effect of direct sorted. I looked into h2 documentation and found very little details. I guess without it the set of records from inner select are sorted before the insert. Since this skip an unneeded step this option should be beneficial.

The change also switched from the iterator looping to indexed for so that's a bonus.

Ready to test/merge to trunk.

LE: is the static inner class CopyMeta of FastCopyHelper useful?

#3 - 11/02/2023 08:36 AM - Greg Shah

- Status changed from Test to WIP

We probably need a different status code for "internal testing before being released to customer". Anyway, let me know when testing is complete and we can get this merged.

#4 - 11/02/2023 08:55 AM - Greg Shah

- Status changed from WIP to Internal Test

Added. :)

#5 - 11/02/2023 09:19 AM - Alexandru Lungu

Wow! This is major :)

I tested with a large customer's regression testing suite and got no problems. I also did some performance tests and got some decent results. The variance is huge (just like stated in #7026-276), but the results are in the current performance bounds.

This is ready for merge.

#6 - 11/02/2023 11:22 AM - Greg Shah

You can merge to trunk.

#7 - 11/02/2023 11:34 AM - Alexandru Lungu

- Status changed from Internal Test to Test

Branch 7993a was merged to trunk as rev. 14814 and archived.

#8 - 11/02/2023 10:52 PM - Ovidiu Maxiniuc

Eugenie reported in #7143-463 an issue related to 7993a.

There are three issues I can see:

- the DDL being executed should be drop table if exists ... instead of drop table if not exists;
- the statements are executed a bit too late, during context cleanup. As such, I wonder whether this extra step is necessary once the temporary tables will be dropped by H2 as the connection is closed;
- before dropping the mappings the user-specific seq_NN is also dropped, but, apparently no error is generated.