

## Wt - Feature #3621

### Db: Consider omitting select count(1) in QueryModel rowCount when possible

10/09/2014 05:31 AM - Bruce Toll

<b>Status:</b>	Closed	<b>Start date:</b>	10/09/2014
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Koen Deforche	<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	3.3.4		

#### Description

When a QueryModel has a batch size large enough to retrieve all rows, it's possible to skip the initial select count(1) used to obtain the row count. This can significantly improve the efficiency of queries where the cost of a SELECT count(1) is high, e.g. when a query includes expensive window functions, *and* the result rows fits within a single batch.

Please see the attached patch for additional documentation. The patch has been lightly tested with the github version of Wt, 3.3.3-19-gad17d0b, and passes the existing suite of Dbo tests, as well as an additional test/benchmark included with the patch. Testing has been limited to Postgres and Sqlite3.

NOTE: One area where performance could potentially degrade is for applications that call rowCount() without subsequently retrieving rows. You may catch other issues upon review. Let me know if you'd like any changes.

#### History

##### #1 - 10/10/2014 08:44 PM - Koen Deforche

- Status changed from *New* to *InProgress*
- Assignee set to *Koen Deforche*
- Target version set to *3.3.4*

##### #2 - 10/16/2014 04:30 PM - Koen Deforche

- Status changed from *InProgress* to *Resolved*

Hey,

Great idea!

To avoid users from suffering a performance hit when only using resultCount(), I've documented that they can avoid this by setting batchSize to 0.

Regards,

koen

##### #3 - 10/22/2014 02:28 AM - Koen Deforche

- Status changed from *Resolved* to *Closed*

#### Files

0001-Reduce-QueryModel-rowCount-SELECTs-for-some-cases.patch 9.53 KB

10/09/2014

Bruce Toll