

Mulgara - Feature #23

We must find a more descriptive way of expressing the cost of iteration than getRowUpperBound()

09/26/2006 11:08 AM - Andrae Muys -

Status:	New	Start date:	
Priority:	Normal	Due date:	
Assignee:	Andrae Muys -	% Done:	0%
Category:	Mulgara	Estimated time:	0.00 hour
Target version:			
Resolution:			

Description

An increasing number of resolvers (relational, federating, distributed) cannot provide a meaningful upper-bound to their row-count.

Also row-count assumes a uniform cost of iteration, when iteration (as opposed to resolution) entails network latency this is unrealistic.

How do we handle this? Do we need some sort of bogomips-like timing constant to provide a baseline for comparison? Can we use this to also provide selectivity data that would also improve our join performance?