

Speeding up pipeline while it's running, OLTP vs. OLAP

In addition to partitioning tables to speed things up, we should really think about adding these tables of active streams and processinstances.

ActiveStream

This new table should be all streams that are not in a finalized state.

ActiveProcessInstance

This is the same. Maybe this makes more sense to do first.

I think we could keep simpler indexes which would keep bookkeeping to a minimum, speeding up transaction time in general and reducing row locks. In addition to this, this might be good if we move to a server-client DB model from the stored procedure model.

I'll need to detail any complications that might arise from this model.