

Database - Feature #1664

data import improvements

10/22/2012 12:33 AM - Eric Faulhaber

Status: New	Start date:
Priority: Normal	Due date:
Assignee:	% Done: 0%
Category:	Estimated time: 0.00 hour
Target version:	vendor_id: GCD
billable: No	
Description	
Related issues:	
Related to Database - Feature #1661: support case-sensitive fields in indexes	Closed 07/19/2013 07/21/2013
Related to Database - Feature #4723: make it significantly easier to run data...	New

History

#1 - 10/22/2012 12:37 AM - Eric Faulhaber

- performance improvement: re-architect threading to distribute work in a more granular way (preserve prioritization by table, but allow multiple threads to run within the import of a single table to avoid the "critical path" problem);
- make schema generation independent of Hibernate (today we rely on Hibernate tools past end-of-life, which have some issues for us);
- we currently cannot re-create the schema to properly use text instead of varchar for string data, due to some changes we made in the P2J<-->Hibernate data bindings;
- make index creation independent of Hibernate (today we let Hibernate generate the index creation DDL; we need more flexibility);
- abstract out PostgreSQL-specific code (running the ANALYZE command after each table is imported) into a more database-independent model.

#2 - 10/31/2012 04:02 PM - Greg Shah

- Target version set to *Deployment and Management Improvements*

#3 - 11/16/2016 01:08 PM - Greg Shah

- Target version deleted (*Deployment and Management Improvements*)

#4 - 12/04/2020 02:14 PM - Eric Faulhaber

- Related to Feature #4723: make it significantly easier to run database import added

#5 - 12/04/2020 02:17 PM - Eric Faulhaber

- Estimated time deleted (48.00)

Note that as of FWD trunk rev 11348, Hibernate is no longer used. As such, the Hibernate-specific issues noted above no longer are relevant.

The other issues remain.