

VegBank - Bug #2402

Strategy to update denormalized data, cache, pages

04/06/2006 05:44 PM - Michael Lee

Status:	New	Start date:	04/06/2006
Priority:	Normal	Due date:	
Assignee:	Michael Lee	% Done:	0%
Category:	performance	Estimated time:	0.00 hour
Target version:	1.2.0	Spent time:	0.00 hour
Bugzilla-Id:	2402		
Description			
We have a series of items that are stored data, views, etc., that help with performance. These items need to be updated at various times when data changes or is added to the database. We need some strategy for how and when to update these denormalizations.			
Related issues:			
Blocks VegBank - Bug #2002: add VegBranch downloads to cache and fix infinite...		Resolved	03/07/2005
Blocks VegBank - Bug #1861: fix denorm SQL approach		Resolved	01/10/2005

History

#1 - 04/06/2006 05:47 PM - Michael Lee

A number of fields are "denorm" fields. These are indicated as attModel='denorm' in /vegbank/docs/xml/db_model_vegbank.xml.

Denorm fields are updated in two ways. The first approach was a .sql file (see /vegbank/src/sql/denorm). Then we added a java utility. I don't know what that's called. The utility uses SQL snippets from the SQLStore.properties file (/vegbank/src/java/org/vegbank/common/SQLStore.properties). We used AJAX to update the denorm fields with this utility or with custom jsp tags (vegbank:denorm). The AJAX approach failed because of http timeouts since the updates take a while sometimes.

#2 - 04/06/2006 05:50 PM - Michael Lee

We also have certain views (stored html which was saved from .jsp files that are dynamic). These are built currently using ant when we build the site (because I didn't know of a better way to do it). This is done by the ant target "cache" in /vegbank/web/build.xml. The views are stored in the "raw" directory, because their contents are just raw, unstyled tables.

#3 - 04/06/2006 05:51 PM - Michael Lee

XML Representations of our plots take some time to write. We haven't yet developed a system to store these in a cache, but we need to (see bug 2002). Once that is developed, we will also have to update the XML cache when plots are updated and/or added.

#4 - 06/01/2006 10:37 AM - Michael Lee

We might be able to use triggers to handle all the denorms. Then it's out of sight and out of mind. Update anything and the triggers handle fixing the denorms to match.

#5 - 08/14/2006 04:41 PM - Michael Lee

I think we decided against triggers as they proved to be slow and potentially circular. This bug remains of high importance. I think it processes correctly after XML Load, but updates should also be triggered after someone interprets a plant or plot. (and that should update the XML_Cache, too)

#6 - 09/11/2006 04:24 PM - Michael Lee

Currently denormalizations are kicked off only when someone loads data. Any subsequent changes to the database need to tie in to the denorm utility and run updates on tables that are affected by the change in question. e.g. when interpreting a plant, you should run denorm update to taxonobservation, which reads some of the denorms from taxoninterpretation, where a new record was just added.

#7 - 03/27/2013 02:20 PM - Redmine Admin

Original Bugzilla ID was 2402