

Morpho - Bug #940

need large performance enhancements

01/10/2003 09:46 AM - Matt Jones

Status:	In Progress	Start date:	01/10/2003
Priority:	Normal	Due date:	
Assignee:	Dan Higgins	% Done:	0%
Category:	morpho - general	Estimated time:	0.00 hour
Target version:	1.9	Spent time:	0.00 hour
Bugzilla-Id:	940		

Description

Morpho is a dog. We need to fix that. There are several obvious places, and probably some more subtle ones. Whoever handles this bug should do some formal application profiling, but here are some suggested places where improvements must be made:

- 1) Opening a data package is v. slow. This is true whether local or network. Note that metacat can open and display a package on the web in about 2 seconds, while Morpho takes 20 seconds or more. We use a ton of XML parsers -- my guess is we need to be much more careful about how things are opened, cache the results, and do some downloading asynchronously based on a priority of what needs to be seen first.
- 2) Searches are slow. Bug 307 proposes one mechanism to improve things which should be pursued. Local searches could easily be improved by changing how docs are stored on disk -- by moving revisions to a separate directory and only searching current docs, we can speed things up a lot. Data should probably be in yet another directory, because it is large, and because it causes an exception to be thrown, which is slower than a normal method call.
- 3) Resultset display is slow. Metacat searches are returned as one large document, rather than being displayed as they are produced. Could possibly change metacat and morpho to immediately stream results rather than holding on to them. Bug 458 describes this approach.

History

#1 - 01/17/2003 03:12 PM - Dan Higgins

I (Dan Higgins) offer the following observations about what is and is not 'slow' in Morpho.

Metacat searches:

- A '%' search for all packages using Morpho takes ~20 sec and returns 169 packages
- A '%' search for all packages using webclient takes ~20 sec and returns 169 packages (no apparent difference from morpho)
- A search for a word in a title which only one package satisfies take 2-3 seconds in both Morpho and using the web client.

Metacat+local searches:

- including 21 local data packages adds less than 1 sec to a % search
- A search for a word in a title which only one package satisfies takes is about the same whether or not the 21 local packages are included in the search (first search is somewhat slower since local cache not yet built)

Opening Data Packages:

Opening packages in Morpho can take 20 secs or more IF there is an associated data file. Packages without a data file open in 2-3 seconds which is about the same time required in the web client. The slowness is apparently tied to the data display. Since the web client only displays datapackage metadata, it would clearly be faster than the Morpho display which shows ALL the package info. One solution for speed up would be to put all the data processing into a background thread and show the package metadata first (while the data thread is running).

I conclude that the primary area where work on speeding up Morpho is in the data

display part of the package display. Local searches are not slowing down the system unless one has hundreds of local packages. Also fixing Bug 307 will have no effect on initial searches although it would speed up Result set refreshes.

#2 - 01/17/2003 03:13 PM - Dan Higgins

I (Dan Higgins)offer the following observations about what is and is not 'slow' in Morpho.

Metacat searches:

--A '%' search for all packages using Morpho takes ~20 sec and returns 169 packages
--A '%' search for all packages using webclient takes ~20 sec and returns 169 packages (no apparent difference from morpho)
--A search for a word in a title which only one package satisfies take 2-3 seconds in both Morpho and using the web client.

Metacat+local searches:

--including 21 local data packages adds less than 1 sec to a % search
--A search for a word in a title which only one package satisfies takes is about the same whether or not the 21 local packages are included in the search (first search is somewhat slower since local cache not yet built)

Opening Data Packages:

Opening packages in Morpho can take 20 secs of more IF there is an associated data file. Packages without a data file open in 2-3 seconds which is about the same time required in the web client. The slowness is apparently tied to the data display. Since the web client only displays datapackage metadata, it would clearly be faster than the Morpho display which shows ALL the package info. One solution for speed up would be to put all the data processing into a background thread and show the package metadata first (while the data thread is running).

I conclude that the primary area where work on speeding up Morpho is in the data display part of the package display. Local searches are not slowing down the system unless one has hundreds of local packages. Also fixing Bug 307 will have no effect on initial searches although it would speed up Result set refreshes.

#3 - 01/17/2003 03:18 PM - Dan Higgins

Didn't mean to close it! (higgins)

#4 - 04/02/2003 12:17 PM - Dan Higgins

A number of corrections have been made to speed up the display of newly opened datapackages. These include removing a large number of redundant XML parsing operations (by caching the DOM) and not trying to parse data files (a list of data files is built so they are newer examined more than once). Also, only the latest version of package is now searched locally.

#5 - 10/07/2003 11:38 AM - Dan Higgins

Newly rewritten XML tree editor (Version 1.4.1) is nearly 100 times faster than older editor. This will make editing og complex modules like 'project' much more reasonable!

#6 - 03/27/2013 02:15 PM - Redmine Admin

Original Bugzilla ID was 940