

Kepler - Bug #5017

ReportLayout LSID and referral list should be maintained

05/19/2010 03:08 PM - Derik Barseghian

Status:	Resolved	Start date:	05/19/2010
Priority:	Normal	Due date:	
Assignee:	Derik Barseghian	% Done:	0%
Category:	reporting	Estimated time:	0.00 hour
Target version:	reporting-2.X.Y	Spent time:	0.00 hour
Bugzilla-Id:	5017		
Description			
When a user changes a reportLayout, its LSID revision should increment. When a user changes a reportLayout with an LSID that has a non-local namespace, the LSID should go into a referral list, and a new LSID should be generated and then incremented on subsequent changes.			

History

#1 - 05/20/2010 11:34 AM - ben leinfelder

are report layouts now NamedObjs or something?
Where is the referral list stored if not?

#2 - 05/20/2010 11:47 AM - Derik Barseghian

They're still not NamedObjs. Unfortunately the constructor for NamedObjIdReferralList needs a NamedObj. Aaron recommended refactoring NamedObjIdReferralList, or having reportLayout keep a NamedObj member to serve as proxy. The latter seems kludgy, but may be the way to go since we're trying to ship 2.0.

One problem currently is that reporting attempts to increment the LSID even if it's from a different namespace, and this is not allowed and thus gives errors. The ideal thing here would be to keep a referral list. One cheat, if we don't care about the origins of a reportLayout, is to simply create a new lsid and discard the old.

#3 - 05/20/2010 11:55 AM - ben leinfelder

Hmm, I'm trying to think of how we'd use the "provenance of the report layout". I suppose that could be useful for undo/redo operations? But it seems like we don't have an immediate need to keep track of history, the more pressing need is to avoid nasty exceptions when trying to update the LSID.

#4 - 05/20/2010 12:12 PM - Derik Barseghian

It might also be nice to see who originally created and who contributed to a report layout, esp in a 'publication kar'. But I agree, discarding may be more prudent at this juncture.

#5 - 06/07/2010 02:50 PM - debi staggs

The report layout LSID is now incremented whenever the user saves a KAR file, but only if it is in the local namespace. If it is not a local instance, we have decided NOT to add it the referral list at this time. Seems to be working well, closing this bug now if there are no objections.

#6 - 06/07/2010 03:42 PM - Derik Barseghian

Hey Debi, while it may be ok to only increment the LSID whenever a report layout is saved, one of the ways a report layout is saved is to provenance when the workflow is run. So if the LSID now only increments during a save to KAR, that won't be sufficient because of the following scenario:

- 1) Design a workflow and report layout
- 2) Execute (call this run1)
- 3) Change the report layout
- 4) Execute (call this run2)
- 5) Export to KAR run1
- 6) Export to KAR run2

run1.kar and run2.kar now contain two different report layouts, but each with the same LSID.

Less important, but just to note: another reason I'm not a fan of this approach is that you can design a workflow and report, save to kar, then save to another kar, and you then have two report layouts that are essentially identical, but simply have different LSIDs (the LSID is the only thing keeping them from being the same).

#7 - 06/07/2010 04:15 PM - debi staggs

So, just to be sure that I am understanding you correctly -- are you saying that using the ReportLayoutKAREntryHandler for this purpose is not

sufficient ? and that we need to increment it when the workflow runs as well ? Can you suggest what you think the best approach to doing that would be ?

Currently, doing it in the ReportLayoutKAREntryHandler using saveReportLayout, does seem to be working very reliably, and I hate to change that.

Thanks.

#8 - 06/08/2010 10:57 AM - Derik Barseghian

Yes, it's not sufficient because scenarios like comment#6 are still possible.

I think the best approach is to roll the report layout's LSID whenever the report layout is changed.

If you don't want to do it this way, but instead only on saves, then you need to roll it before any save, which includes before the save to provenance when the workflow is run.

#9 - 06/09/2010 12:38 PM - debi staggs

(In reply to comment [#8](#))

Yes, it's not sufficient because scenarios like comment#6 are still possible.

I think the best approach is to roll the report layout's LSID whenever the report layout is changed.

If you don't want to do it this way, but instead only on saves, then you need to roll it before any save, which includes before the save to provenance when the workflow is run.

O.K. sounds like a good idea. I have now had it update the revision whenever files are written to Provenance as well. Does that seem like it will adequately address the issues you have described ?

#10 - 06/09/2010 01:01 PM - debi staggs

(In reply to comment [#9](#))

(In reply to comment [#8](#))

Yes, it's not sufficient because scenarios like comment#6 are still possible.

I think the best approach is to roll the report layout's LSID whenever the report layout is changed.

If you don't want to do it this way, but instead only on saves, then you need to roll it before any save, which includes before the save to provenance when the workflow is run.

O.K. sounds like a good idea. I have now had it update the revision whenever files are written to Provenance as well. Does that seem like it will adequately address the issues you have described ?

Also I have now tested this using the same steps described in your Comment [#6](#).

Upon viewing the ROML LSIDS for each run, they are now different, as well as incremented sequentially by 1, each time the workflow is run, AND/OR each time a KAR is saved.

Please review and let me know if you agree, so that we can close this bug for now.

Thanks.

#11 - 06/09/2010 01:10 PM - Derik Barseghian

Sounds good. Can you poke around looking for any other similar problems -- e.g. what happens when you have multiple copies of the same report open (same LSID), make changes to one, but switch to the original window and execute that workflow? There exist in the system at that moment two report layouts with the same LSID. Hopefully the right report is used and has its LSID incremented, and hopefully the other report doesn't lose its changes, will have its LSID incremented and be saved to prov when that wf is run, etc.

#12 - 06/14/2010 12:21 PM - debi staggs

(In reply to comment [#11](#))

Sounds good. Can you poke around looking for any other similar problems -- e.g. what happens when you have multiple copies of the same report open (same LSID), make changes to one, but switch to the original window and execute that

workflow? There exist in the system at that moment two report layouts with the same LSID. Hopefully the right report is used and has its LSID incremented, and hopefully the other report doesn't lose its changes, will have its LSID incremented and be saved to prov when that wf is run, etc.

After testing I have found that what happens in this case, is that when you open two copies of the same workflow, select one of them, make changes to one, then switch to the original window and execute that workflow, then the LSID for the report layout is incremented at that point. There are no longer 'two' report layouts. There is one current layout, and one layout that has (in essence) been changed (by running the workflow) but not saved. If you switch to the workflow in either window they will both show the same LSID for the report layout. If you go to that window where the first change was made it will always reload the layout from the workflow no matter what. This is something that we chose not to address for this release and can be referenced in Bug [#5009](#) (i.e. Opening two KARS that have the same Workflow, but that have different report layouts, will make the first layout refresh itself with the layout from the KAR that was opened most recently.) which still applies.

Can I suggest that this is satisfactory for now and that we should close this bug ?

#13 - 06/14/2010 01:12 PM - Derik Barseghian

(In reply to comment [#12](#))

(In reply to comment [#11](#))

Sounds good. Can you poke around looking for any other similar problems -- e.g. what happens when you have multiple copies of the same report open (same LSID), make changes to one, but switch to the original window and execute that workflow? There exist in the system at that moment two report layouts with the same LSID. Hopefully the right report is used

I gave the above a try, and what I found (by looking at the generated report) is that the wrong layout is used (the other window's layout--the one modified by user-- is used).

and has its LSID incremented, and hopefully the other report doesn't lose its changes, will have its LSID incremented and be saved to prov when that wf is run, etc.

After testing I have found that what happens in this case, is that when you open two copies of the same workflow, select one of them, make changes to one, then switch to the original window and execute that workflow, then the LSID for the report layout is incremented at that point. There are no longer 'two' report layouts. There is one current layout, and one layout that has (in essence) been changed (by running the workflow) but not saved. If you switch to the workflow in either window they will both show the same LSID for the report layout.

How are you looking at the report layout here? I'm confused.

If you go to that window where the first change was made it will always reload the layout from the workflow no matter what.

Confused here also. For the above scenario, my unmodified layout stays unmodified in the gui. This bug would be a little less confusing if the layout did refresh to show the changes in the other window, since they will be used during execution.

This is something that we chose not to address for this release and can be referenced in Bug [#5009](#) (i.e. Opening two KARS that have the same Workflow, but that have different report layouts, will make the first layout refresh itself with the layout from the KAR that was opened most recently.) which still applies.

Can I suggest that this is satisfactory for now and that we should close this bug ?

I wouldn't close it because I'm guessing incrementing the lsid whenever the report layout is changed is a good way to solve the bug. If we want to add a caveat about not being able to have two copies of the same workflow+report layout pair opened at once then we could maybe postpone the bug.

#14 - 06/14/2010 01:39 PM - debi staggs

(In reply to comment [#13](#))

(In reply to comment [#12](#))

(In reply to comment [#11](#))

Sounds good. Can you poke around looking for any other similar problems -- e.g. what happens when you have multiple copies of the same report open (same LSID), make changes to one, but switch to the original window and execute that workflow? There exist in the system at that moment two report layouts with the same LSID. Hopefully the right report is used

I gave the above a try, and what I found (by looking at the generated report) is that the wrong layout is used (the other window's layout--the one modified by user-- is used).

and has its LSID incremented, and hopefully the other report doesn't lose its changes, will have its LSID incremented and be saved to prov when that wf is run, etc.

After testing I have found that what happens in this case, is that when you open two copies of the same workflow, select one of them, make changes to one, then switch to the original window and execute that workflow, then the LSID for the report layout is incremented at that point. There are no longer 'two' report layouts. There is one current layout, and one layout that has (in essence) been changed (by running the workflow) but not saved. If you switch to the workflow in either window they will both show the same LSID for the report layout.

How are you looking at the report layout here? I'm confused.

If you go to that window where the first change was made it will always reload the layout from the workflow no matter what.

Confused here also. For the above scenario, my unmodified layout stays unmodified in the gui. This bug would be a little less confusing if the layout did refresh to show the changes in the other window, since they will be used during execution.

This is something that we chose not to address for this release and can be referenced in Bug [#5009](#) (i.e. Opening two KARS that have the same Workflow, but that have different report layouts, will make the first layout refresh itself with the layout from the KAR that was opened most recently.) which still applies.

Can I suggest that this is satisfactory for now and that we should close this bug ?

I wouldn't close it because I'm guessing incrementing the lsid whenever the report layout is changed is a good way to solve the bug. If we want to add a caveat about not being able to have two copies of the same workflow+report layout pair opened at once then we could maybe postpone the bug.

To document our conversation from the REAP meeting today, it sounds like the more basic problem to all of this is the fact that the KAR only references the last report layout that is generated. We should definitely address this as soon as possible. However, since it will take a significant effort and will have a fairly wide impact - it may need to be delayed until the next release.

#15 - 06/21/2010 03:06 PM - debi staggs

(In reply to comment [#14](#))

(In reply to comment [#13](#))

(In reply to comment [#12](#))

(In reply to comment [#11](#))

Sounds good. Can you poke around looking for any other similar problems -- e.g. what happens when you have multiple copies of the same report open (same LSID), make changes to one, but switch to the original window and execute that workflow? There exist in the system at that moment two report layouts with the same LSID. Hopefully the right report is used

I gave the above a try, and what I found (by looking at the generated report) is that the wrong layout is used (the other window's layout--the one modified by user-- is used).

and has its LSID incremented, and hopefully the other report doesn't lose its changes, will have its LSID incremented and be saved to prov when that wf is run, etc.

After testing I have found that what happens in this case, is that when you open two copies of the same workflow, select one of them, make changes to one, then switch to the original window and execute that workflow, then the LSID for the report layout is incremented at that point. There are no longer 'two' report layouts. There is one current layout, and one layout that has (in essence) been changed (by running the workflow) but not saved. If you switch to the workflow in either window they will both show the same LSID for the report layout.

How are you looking at the report layout here? I'm confused.

If you go to that window where the first change was made it will always reload the layout from the workflow no matter what.

Confused here also. For the above scenario, my unmodified layout stays unmodified in the gui. This bug would be a little less confusing if the layout did refresh to show the changes in the other window, since they will be used during execution.

This is something that we chose not to address for this release and can be referenced in Bug [#5009](#) (i.e. Opening two KARS that have the same Workflow, but that have different report layouts, will make the first layout refresh itself with the layout from the KAR that was opened most recently.) which still applies.

Can I suggest that this is satisfactory for now and that we should close this bug ?

I wouldn't close it because I'm guessing incrementing the lsid whenever the report layout is changed is a good way to solve the bug. If we want to add a caveat about not being able to have two copies of the same workflow+report layout pair opened at once then we could maybe postpone the bug.

To document our conversation from the REAP meeting today, it sounds like the more basic problem to all of this is the fact that the KAR only references the last report layout that is generated. We should definitely address this as soon as possible. However, since it will take a significant effort and will have a fairly wide impact - it may need to be delayed until the next release.

After the REAP meeting today, Matt, Derik, and I discussed some of the issues related to this bug. In the course of our discussion the following problem was described:

```
Window 1
Workflow [A]      KAR 0
Layout [0]

Window 2
Workflow [A]      KAR0 -> KAR1
Layout [0] -> Layout1

Workflow [B]
  Layout [1]
  Layout [2]
  Layout [3]
```

The following solutions were proposed:

1. LayoutChangedEvent to sync layouts on change via refresh
-- might be hard, and isn't the right direction

2. Support multiple layouts per wf
-- a. gui menu to save layout, menu to switch in RD, menu to switch in report viewer, render multiple reports.
-- b. could just allow layouts to diverge, association maintained in manager that uses tableau window and layout ID (rather than Workflow LSID)

3. Don't allow the same wf to be opened more than once
-- would limit the ability of WRM to open multiple runs of same wf.

After some discussion, we decided that solution 2b would provide the best functionality at this time, and would result in application behaviors that users could reasonably anticipate.

#16 - 06/30/2010 06:16 PM - Derik Barseghian

As of r25050 I believe I have the report LSID being properly maintained - if the user deletes, adds, rearranges, or edits report items, or changes their properties, or changes the title, the LSID will roll. This could be improved, but I think it's ok as is.

(Improvements:

- all text areas use a DocumentListener for changes, so any letter you type or delete rolls the LSID - if the GUI required the user do something to 'complete' their changes to any given text area, I could then do a comparison between the old and new values and only roll the LSID once.
 - the image report item will roll LSID whenever the user clicks on the '+ Get Image' button - it would be better to compare new and previous Images to see if they are different before rolling LSID
 - when dragging a panel and dropping it back in the location from which you started will roll the LSID - no need to roll here
-)

We decided to not maintain the referral list, so if the report has a non local lsid it gets tossed and a new one is used.

These changes are in 2.0 only for now, I'll merge into trunk later.

Now the issue in comment#15 and bug#5009 can probably be addressed.

#17 - 07/06/2010 02:07 PM - Derik Barseghian

Comment#15 should be fixed now in 2.0 at r25056. Needs to be merged into trunk.

#18 - 02/01/2011 03:06 PM - Derik Barseghian

I'm having trouble understanding the "problem" described in comment#15 at this point, so I'll just trust my comment#17 saying I fixed comment#15. I merged everything a long time ago, so closing.

#19 - 03/27/2013 02:28 PM - Redmine Admin

Original Bugzilla ID was 5017