

Arvados - Bug #13006

is_a filter doesn't work with remote objects

01/26/2018 04:21 PM - Lucas Di Pentima

Status:	Resolved	Start date:	12/13/2018
Priority:	Normal	Due date:	
Assigned To:	Lucas Di Pentima	% Done:	100%
Category:	API	Estimated time:	0.00 hour
Target version:	2018-12-21 Sprint		
Description			
While working on #12758 I found that the sync-groups tool was failing to retrieve group membership links from remote (federated) user accounts. Seems that we need to modify the API server to make a substring query on tail_uuid & head_uuid fields, so that the uuid prefix is not fixed.			
Subtasks:			
Task # 14610: Review 13006-api-is_a-filter & 13006-sync-groups-is_a-filter			Resolved

Associated revisions

Revision 55c142e4 - 12/14/2018 04:02 PM - Lucas Di Pentima

Merge branch '13006-api-is_a-filter'
Refs #13006

Arvados-DCO-1.1-Signed-off-by: Lucas Di Pentima <ldipentima@veritasgenetics.com>

Revision cd4bb0b5 - 12/14/2018 04:03 PM - Lucas Di Pentima

Merge branch '13006-sync-groups-is_a-filter'
Closes #13006

Arvados-DCO-1.1-Signed-off-by: Lucas Di Pentima <ldipentima@veritasgenetics.com>

History

#2 - 01/26/2018 04:53 PM - Peter Amstutz

This is implemented in arvados/services/api/lib

```
when 'is_a'
  operand = [operand] unless operand.is_a? Array
  cond = []
  operand.each do |op|
    cl = ArvadosModel::kind_class op
    if cl
      cond << "#{ar_table_name}.#{attr} like ?"
      param_out << cl.uuid_like_pattern
    else
      cond << "1=0"
    end
  end
  cond_out << cond.join(' OR ')
```

Instead of using 'like' this should use a substring match:

```
operand.each do |op|
  cl = ArvadosModel::kind_class op
  if cl
    cond << "substring(#{ar_table_name}.#{attr}, 6, 5) = ?"
    param_out << cl.uuid_prefix
  else
    cond << "1=0"
  end
end
```

If 'attr' is 'uuid' I think you can even short circuit the check with "model_class.uuid_prefix == cl.uuid_prefix" (we can assume every value in the uuid column has the same type.)

To ensure the substring match is efficient, we may need to create an expression index:

<https://www.postgresql.org/docs/9.4/static/indexes-expressional.html>

#3 - 02/01/2018 08:36 PM - Lucas Di Pentima

- Assigned To deleted (Lucas Di Pentima)

#4 - 02/01/2018 09:07 PM - Tom Morris

- Target version set to To Be Groomed

#5 - 02/07/2018 06:12 PM - Tom Morris

- Story points set to 1.0

#6 - 04/18/2018 03:42 PM - Tom Morris

- Target version changed from To Be Groomed to Arvados Future Sprints

#7 - 12/12/2018 05:09 PM - Tom Morris

- Target version changed from Arvados Future Sprints to 2018-12-21 Sprint

#8 - 12/12/2018 07:42 PM - Lucas Di Pentima

- Status changed from New to In Progress

#9 - 12/13/2018 04:38 PM - Lucas Di Pentima

- Assigned To set to Lucas Di Pentima

#10 - 12/13/2018 07:26 PM - Lucas Di Pentima

Updates at [cd7a746df](#) - branch 13006-api-is_a-filter

Test run: <https://ci.curoverse.com/job/developer-run-tests/1009/>

- Followed Peter's suggestion by changing the query to use substring(), avoiding the check when attr is "uuid" and prefix_uuid matches.
- Added expression index on links.[tail|head]_uuid
- Added test

I also could update the group-sync tool to use the is_a operator again, for enhanced performance.

#11 - 12/13/2018 08:04 PM - Peter Amstutz

Lucas Di Pentima wrote:

Updates at [cd7a746df](#) - branch 13006-api-is_a-filter

Test run: <https://ci.curoverse.com/job/developer-run-tests/1009/>

- Followed Peter's suggestion by changing the query to use substring(), avoiding the check when attr is "uuid" and prefix_uuid matches.

If attr == 'uuid' but model_class.uuid_prefix != cl.uuid_prefix, should it fail unconditionally?

- Added expression index on links.[tail|head]_uuid
- Added test

I also could update the group-sync tool to use the is_a operator again, for enhanced performance.

How about in a separate branch.

#12 - 12/13/2018 08:57 PM - Lucas Di Pentima

Peter Amstutz wrote:

If attr == 'uuid' but model_class.uuid_prefix != cl.uuid_prefix, should it fail unconditionally?

I believe so, fixed.

Update at [96165e8fc](#)

Test run: <https://ci.curoverse.com/job/developer-run-tests/1011/>

#13 - 12/13/2018 09:05 PM - Lucas Di Pentima

Peter Amstutz wrote:

I also could update the group-sync tool to use the is_a operator again, for enhanced performance.

How about in a separate branch.

Done: [84dcca64f](#) - branch 13006-sync-groups-is_a-filter

Test run: <https://ci.curoverse.com/job/developer-run-tests/1012/> (not yet started atm)

#14 - 12/14/2018 03:22 PM - Peter Amstutz

Lucas Di Pentima wrote:

Peter Amstutz wrote:

I also could update the group-sync tool to use the is_a operator again, for enhanced performance.

How about in a separate branch.

Done: [84dcca64f](#) - branch 13006-sync-groups-is_a-filter

Test run: <https://ci.curoverse.com/job/developer-run-tests/1012/> (not yet started atm)

Is there a 'sync-groups' test that fails without the '13006-api-is_a-filter' fix but passes with it?

#15 - 12/14/2018 03:28 PM - Lucas Di Pentima

Yes. Although I didn't tested it now, but it was the means to detect this is_a bug in the API server. The test is TestMembershipRemoval, if you want I can rebase this branch to master and then merge the other 13006 branch to prove that the test passes.

#16 - 12/14/2018 03:32 PM - Peter Amstutz

Lucas Di Pentima wrote:

Yes. Although I didn't tested it now, but it was the means to detect this is_a bug in the API server. The test is TestMembershipRemoval, if you want I can rebase this branch to master and then merge the other 13006 branch to prove that the test passes.

Yes, I think that would be a good idea. Go ahead and merge both branches if it goes as expected (test fails without fix → test passes with fix)

#17 - 12/14/2018 03:57 PM - Lucas Di Pentima

Ok, so I rebased 13006-sync-groups-is_a-filter ([0c21d2bd2d0fd351a4e546002a249b0b748061eb](#)) to master. Ran the sync-groups tests locally and got the failure:

```
FAIL: sync-groups_test.go:267: TestSuite.TestMembershipRemoval
```

```
sync-groups_test.go:308:
  c.Assert(GroupMembershipExists(s.cfg.Client, remoteUserUUID, groupUUID), Equals, false)
... obtained bool = true
... expected bool = false
```

Then, merged 13006-api-is_a-filter into 13006-sync-groups-is_a-filter ([cbf93e8d897448dbd52369afe89fef2392140ff1](#)). Tests now pass.

#18 - 12/14/2018 04:17 PM - Lucas Di Pentima

- Status changed from In Progress to Resolved

- % Done changed from 0 to 100

Applied in changeset [arvados|cd4bb0b5bde62d7fe32cc105d6311dd318ebe304](#).

#19 - 03/01/2019 06:30 PM - Tom Morris

- Release set to 15