

## Arvados - Bug #11261

### [API] full text search on pipeline\_instances table is extremely slow (wrong index?)

03/15/2017 07:29 PM - Tom Clegg

<b>Status:</b>	New	<b>Start date:</b>	
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>		<b>% Done:</b>	0%
<b>Category:</b>	Performance	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>			
<b>Description</b>			
<b>Related issues:</b>			
Copied from Arvados - Bug #11071: [Workbench] qr1hi performance tests are broken		<b>Resolved</b>	<b>03/14/2017</b>

### History

#### #1 - 03/15/2017 07:30 PM - Tom Clegg

Copied from [#11071](#):

Version query:

```
curl https://${ARVADOS_API_HOST}/discovery/v1/apis/arvados/v1/rest | jq .source_version
```

Test query:

```
time curl -H "Authorization: OAuth2 $ARVADOS_API_TOKEN" -d count=none -d filters='[["any", "@@", "hash:*"]]' -d limit=1 -d _method=GET https://${ARVADOS_API_HOST}/arvados/v1/groups/contents | jq '.items[0].href'
```

results:

version	cluster	count=	real
7eda588-7	qr1hi	none	1m19.801s
7eda588-7	qr1hi	exact	1m20.374s
b1aa6c8-7	4xphq	none	0m0.046s
b1aa6c8-7	4xphq	exact	0m7.106s

Another test query that ensures we query the collections table:

```
time curl -H "Authorization: OAuth2 $ARVADOS_API_TOKEN" -d count=none -d filters='[["any", "@@", "hash:*"], ["uuid", "is_a", "arvados#collection"]]' -d limit=1 -d _method=GET https://${ARVADOS_API_HOST}/arvados/v1/groups/contents | jq '.items[0].href'
```

results:

version	cluster	count=	real
7eda588-7	qr1hi	none	0m0.134s
7eda588-7	qr1hi	exact	0m0.143s
b1aa6c8-7	4xphq	none	0m0.051s
b1aa6c8-7	4xphq	exact	0m0.063s

Hm. Tried a few other queries on qr1hi, and all of the slowness seems to be the fault of the pipeline\_instances table:

count=	uuid is_a ...	real
exact	arvados#pipelineTemplate	0m0.059s
exact	arvados#containerRequest	0m0.186s
exact	arvados#job	0m0.106s
exact	arvados#workflow	0m0.119s

exact	arvados#group	0m0.068s
exact	arvados#collection	0m0.143s
exact	arvados#pipelineInstance	2m40.890s
none	arvados#pipelineInstance	2m41.996s

## #2 - 03/15/2017 07:31 PM - Tom Clegg

```
tomclegg@shell.qlhi:~$ time curl -H "Authorization: OAuth2 $ARVADOS_API_TOKEN" -d count=none -d filters='[["any", "@@", "hash:*"]]' -d limit=1 -d _method=GET https://${ARVADOS_API_HOST}/arvados/v1/pipeline_instances | jq '.items[0].href'
```

...

```
real    2m41.471s
user    0m0.028s
sys     0m0.048s
```

## #3 - 05/10/2017 06:07 AM - Tom Morris

- Category set to Performance
- Target version set to 2017-06-07 sprint

## #4 - 05/24/2017 07:05 PM - Tom Morris

- Target version changed from 2017-06-07 sprint to Arvados Future Sprints

## #5 - 07/06/2021 09:24 PM - Peter Amstutz

- Target version deleted (Arvados Future Sprints)