

## D4Science Infrastructure - Task #9368

### JackRabbit Repo cluster over Postgres for development infra

Jul 24, 2017 12:20 PM - Massimiliano Assante

|   |   |                        |              |
|---|---|------------------------|--------------|
| <b>Status:</b>  | Closed  | <b>Start date:</b>     | Jul 24, 2017 |
| <b>Priority:</b>  | High  | <b>Due date:</b>       | Aug 04, 2017 |
| <b>Assignee:</b>  | _InfraScience Systems Engineer                  | <b>% Done:</b>         | 100%         |
| <b>Category:</b>  | Other   | <b>Estimated time:</b> | 0.00 hour    |
| <b>Target version:</b>  | Moving of Dev and preprod DB to shared Postgres |                        |              |
| <b>Infrastructure:</b>  | Development                                     |                        |              |
| <b>Description</b><br>The current JackRabbit Repository instance for development is going to be dismissed and will be replaced by the same setting we have in production. With the only difference that the cluster is going to be fully operative.<br><br>We have 2 nodes already available:<br><br>workspace-repository1-d.d4science.org<br>workspace-repository2-d.d4science.org<br><br>And we are going to use postgresql-srv-dev.d4science.org for the DB. |   |                        |              |

#### History

##### #2 - Jul 24, 2017 12:21 PM - Massimiliano Assante

- Description updated

##### #3 - Jul 25, 2017 07:19 PM - Andrea Dell'Amico

- Infrastructure deleted (Pre-Production)

I need the correct scopes

##### #4 - Jul 25, 2017 08:39 PM - Andrea Dell'Amico

Another piece: the version I have in the playbook is 1.8.0-4.6.0-149356 but I'm sure that a newer version was manually installed. What's the latest good one?

##### #5 - Jul 26, 2017 05:36 PM - Andrea Dell'Amico

- Status changed from New to In Progress

The hostname configured on the load balancer is workspace-repository.dev.d4science.org.

##### #6 - Jul 26, 2017 05:40 PM - Costantino Perciante

As far as the list of scopes

/gcube  
/gcube/devsec  
/gcube/devNext

but calls coming from vres must be accepted as well (authorizedChildren property set to true in container.xml)

##### #7 - Jul 26, 2017 08:06 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

- % Done changed from 0 to 90

Both servers have been provisioned.

The paths to the jackrabbit data are identical to the production one, with /home/gcube/tomcat/jacrabbit that points to /data/jackrabbit (no dedicated file system, this time).

I stopped tomcat on workspace-repository2-d.d4science.org (workspace-repository-devsec.d4science.org) so that there's a chance to test the cluster setup.

About the cluster: I don't know if the cluster ID must be numeric: I set it to the VM hostname.  
I also already set the janitorEnabled value to true on both servers.

Remember to not start the second server before adding the required cluster information into the database.

**#8 - Jul 27, 2017 09:34 AM - Costantino Perciante**

- % Done changed from 90 to 0

Andrea Dell'Amico wrote:

Both servers have been provisioned.

The paths to the jackrabbit data are identical to the production one, with /home/gcube/tomcat/jacrabbit that points to /data/jackrabbit (no dedicated file system, this time).

I stopped tomcat on workspace-repository2-d.d4science.org (workspace-repository-devsec.d4science.org) so that there's a chance to test the cluster setup.

About the cluster: I don't know if the cluster ID must be numeric: I set it to the VM hostname.  
I also already set the janitorEnabled value to true on both servers.

Remember to not start the second server before adding the required cluster information into the database.

Thanks Andrea, we will report here the results after some tests

**#9 - Jul 27, 2017 11:37 AM - Costantino Perciante**

- % Done changed from 0 to 90

**#10 - Aug 03, 2017 10:10 AM - Costantino Perciante**

Yesterday we take some hours to add the second node to the cluster and start up it as well, make some tests and so on. The configuration seems ok, and there was no problem in adding it despite the fact that the cleaner thread was active since the beginning. Another test I did was to check how the cleaner behaves when one of the node is down (but still belongs to the cluster): it just cleans up the revision information no longer needed by both nodes (that's reasonable).

However, I've removed again workspace-repository-devsec.d4science.org from the cluster, since we would like to have the storage cleaned before we go further with tests.

P.S.

to remove a node from the cluster, its node\_id must be removed from local\_revisions table

**#11 - Aug 13, 2017 03:54 PM - Andrea Dell'Amico**

I've added a wiki page with the steps needed to add/remove a node to/from the cluster, here:

[https://support.d4science.org/projects/aginfraplut/wiki/JackRabbit\\_\(Home\\_Library\)](https://support.d4science.org/projects/aginfraplut/wiki/JackRabbit_(Home_Library))

Feel free to amend it if there's something wrong or that's not clear.

**#12 - Sep 13, 2017 03:18 PM - Costantino Perciante**

Andrea Dell'Amico wrote:

I've added a wiki page with the steps needed to add/remove a node to/from the cluster, here:

[https://support.d4science.org/projects/aginfraplut/wiki/JackRabbit\\_\(Home\\_Library\)](https://support.d4science.org/projects/aginfraplut/wiki/JackRabbit_(Home_Library))

Feel free to amend it if there's something wrong or that's not clear.

Thanks Andrea, it sounds ok to me.

**#13 - Nov 20, 2017 04:04 PM - Andrea Dell'Amico**

The infrastructure is in place since months. The cluster features aren't used yet but it's not a matter of configuration.

**#14 - Nov 20, 2017 05:57 PM - Andrea Dell'Amico**

- Status changed from Feedback to Closed

- % Done changed from 90 to 100