

D4Science Infrastructure - Task #3475

Add 2 OrientDB instances

Apr 15, 2016 06:18 PM - Luca Frosini

Status:	Closed	Start date:	Sep 29, 2016
Priority:	Normal	Due date:	
Assignee:	Andrea Dell'Amico	% Done:	100%
Category:	Application	Estimated time:	0.00 hour
Target version:	Dev Resource Registry Deployment		
Infrastructure:	Development		
Description			
Subtasks:			
Task # 5227: Enable SSL on the orientdb instances			Closed
Related issues:			
Related to D4Science Infrastructure - Task #4984: Create new VM to host new R...		Closed	Sep 09, 2016 Sep 13, 2016
Related to D4Science Infrastructure - Task #5568: Please Update orientdb vers...		Closed	Oct 26, 2016

History

#1 - Apr 15, 2016 06:19 PM - Luca Frosini

- Parent task set to #3474

#2 - Apr 15, 2016 06:19 PM - Luca Frosini

- Tracker changed from Support to Task

#3 - Apr 15, 2016 06:22 PM - Luca Frosini

- Start date changed from Apr 15, 2016 to Mar 07, 2016

#4 - May 06, 2016 05:48 PM - Luca Frosini

- Due date set to Mar 11, 2016

I'll provide configuration information in the next days

#5 - May 10, 2016 03:03 PM - Luca Frosini

- Subject changed from Add other 2 OrientDB instances to Add 2 OrientDB instances

There are 3 file to be configured through ansible scripts:

- ~/orientdb-community-2.1.13/config/orientdb-server-config.xml
- ~/orientdb-community-2.1.13/config/hazelcast.xml
- ~/orientdb-community-2.1.13/config/default-distributed-db-config.json

I propose to start in the following way:

I would like test the cluster so that adding a new server instance must not imply to restart the others (especially on production). To accomplish this we can configure hazelcast to use:

- multicast
- encryption on server synchronizations.

Due to actual sharding limitations (<http://orientdb.com/docs/last/Distributed-Sharding.html#limitation>) I would like to avoid it (seems that they will improve this part in the next version 2.2).

In particular:

Regarding **orientdb-server-config.xml**:

- it is important to force the distributed mode as explained here <http://orientdb.com/docs/last/Distributed-Configuration.html#orientdbserverconfigxml>

```
<parameter name="enabled" value="true" />
```

- You can decide if you want to store the admin password as plain text or as hash:

```
<user resources="*" password="7EA0EAF7FCCECA5C5B261B69CD8EC0F725FB8B3102AF51A4033CC3DED22EBD8A" name="root"/>
<user resources="*" password="PlainPassword" name="root"/>
```

- Please note that during the development I will have the need to modify this configuration by adding tag. An example can be:

```
<hooks>
  <hook class="org.gcube.informationssystem.hooks.DocumentHook" position="REGULAR"/>
  <hook class="org.gcube.informationssystem.hooks.RecordHook" position="REGULAR"/>
</hooks>
```

Regarding **hazelcast.xml**:

- please use multicast configuration <http://orientdb.com/docs/last/Distributed-Architecture-Lifecycle.html#auto-discovering>
- Use security (you can select PBEWithMD5AndDES or what you think is more secure) can find detailed information at the following links: <http://orientdb.com/docs/last/Distributed-Architecture-Lifecycle.html#security>

Regarding **default-distributed-db-config.json** please remove **clusters** property.

More details can be found here:

<http://orientdb.com/docs/last/Distributed-Configuration.html>

Feel free to contact me for any doubt.

#6 - May 11, 2016 12:32 PM - Luca Frosini

Can you also create a txt DNS record as described here:

<http://orientdb.com/docs/last/Distributed-Configuration.html#use-the-dns>

#7 - May 11, 2016 12:32 PM - Luca Frosini

- Priority changed from Normal to High

#8 - May 11, 2016 12:38 PM - Luca Frosini

- Parent task changed from #3474 to #3897

#10 - May 17, 2016 04:29 PM - Luca Frosini

- Priority changed from High to Urgent

#11 - May 25, 2016 01:36 PM - Tommaso Piccioli

- Assignee changed from _InfraScience Systems Engineer to Luca Frosini

- % Done changed from 0 to 60

orientdb01-d-d4s.d4science.org 146.48.122.23
orientdb02-d-d4s.d4science.org 146.48.122.14
orientdb03-d-d4s.d4science.org 146.48.122.148

No ansible at the moment for the orientdb/hazelcast configuration, it will be done later, please configure the cluster by hand if you want to test it.

You can connect via ssh as orientdb user and manage orientdb service with
/home/orientdb/orientdb start/stop/status

#12 - May 25, 2016 02:13 PM - Luca Frosini

- Status changed from New to In Progress

Ok. I'll take care of configuring the cluster.

#13 - May 25, 2016 05:53 PM - Luca Frosini

- Assignee changed from Luca Frosini to Tommaso Piccioli

@tommaso.piccioli@isti.cnr.it I need your support to correctly configure hazelcast.

#14 - May 31, 2016 03:40 PM - Luca Frosini

- Assignee changed from Tommaso Piccioli to _InfraScience Systems Engineer

Can you configure hazelcast with static hostnames? We will try multicast in the future.

#15 - Jun 28, 2016 10:56 AM - Luca Frosini

- % Done changed from 60 to 100

#16 - Jun 28, 2016 10:56 AM - Luca Frosini

- Status changed from In Progress to Closed

#17 - Sep 09, 2016 10:59 AM - Luca Frosini

- Status changed from Closed to In Progress

- Priority changed from Urgent to Normal

- Target version changed from OrientDB Assessment to Dev Resource Registry Deployment

- % Done changed from 100 to 70

We need to setup OrientDB cluster with the 3 VMs available

#18 - Sep 19, 2016 05:22 PM - Andrea Dell'Amico

- Related to Task #4984: Create new VM to host new Resource Registry added

#19 - Sep 23, 2016 07:10 PM - Andrea Dell'Amico

- Assignee changed from _InfraScience Systems Engineer to Andrea Dell'Amico

- % Done changed from 70 to 50

I'm assembling a orientdb ansible role to provision all the pieces.

#20 - Sep 28, 2016 07:10 PM - Andrea Dell'Amico

- % Done changed from 50 to 90

The playbook is ready and the configuration seems corret. We have errors on the hooks, because we didn't create the rules that install the needed jars. I need directions for those.

#21 - Sep 28, 2016 07:13 PM - Andrea Dell'Amico

We also need to tweak the java heap configuration.

#22 - Sep 29, 2016 12:19 PM - Luca Frosini

Andrea Dell'Amico wrote:

The playbook is ready and the configuration seems corret. We have errors on the hooks, because we didn't create the rules that install the needed jars. I need directions for those.

The jar file is available here:

<http://maven.research-infrastructures.eu/nexus/service/local/artifact/maven/redirect?r=gcube-snapshots&g=org.gcube.information-system&a=resource-registry-orientdb-hooks&v=1.0.0-SNAPSHOT&e=jar&c=jar-with-dependencies>

Maven coordinates:

```
<groupId>org.gcube.information-system</groupId>
<artifactId>resource-registry-orientdb-hooks</artifactId>
<version>1.0.0-SNAPSHOT</version>
```

The jar must be copied in \$ORIENTDB_HOME/lib. Every time the lib is changed orient must be restarted.

Class to be included in **orientdb-server-config.xml** is **org.gcube.informationssystem.orientdb.hooks.HeaderHook**.
The result in **orientdb-server-config.xml** must be :

```
<hooks>
  <hook class="org.gcube.informationssystem.orientdb.hooks.HeaderHook" position="REGULAR"/>
</hooks>
```

#23 - Sep 29, 2016 12:56 PM - Andrea Dell'Amico

- *Status changed from In Progress to Feedback*
- *% Done changed from 90 to 100*

Done. The hook has been installed, the configuration updated, the services restarted.

#24 - Sep 29, 2016 03:41 PM - Andrea Dell'Amico

The java heap is now configurable. I've set it to 2048MB.

#25 - Oct 26, 2016 11:42 AM - Luca Frosini

- *Status changed from Feedback to Closed*

#26 - Oct 26, 2016 11:42 AM - Luca Frosini

- *Status changed from Closed to In Progress*

#27 - Oct 26, 2016 11:47 AM - Luca Frosini

- *Related to Task #5568: Please Update orientdb version to 2.2.12 added*

#28 - Apr 18, 2017 05:38 PM - Andrea Dell'Amico

- *Status changed from In Progress to Closed*