

D4Science Infrastructure - Task #140

36 new smartgears VMs are needed

May 23, 2015 02:09 PM - Andrea Dell'Amico

Status:	Closed	Start date:	May 23, 2015
Priority:	High	Due date:	
Assignee:	Andrea Dell'Amico	% Done:	100%
Category:	System Application	Estimated time:	0.00 hour
Target version:	System Configuration		
Infrastructure:	Production		
Description			
Gianpaolo needs 36 more smartgears worker nodes. They need both the latest smartgears software and R installed as described on task #131 .			
Related issues:			
Related to D4Science Infrastructure - Task #131: Report the R Interpreter env...		Closed	May 22, 2015
Related to D4Science Infrastructure - Task #138: Automate the installation of...		Closed	May 22, 2015
Related to D4Science Infrastructure - Task #139: Automate the smartgears inst...		Closed	May 23, 2015

History

#1 - May 23, 2015 02:10 PM - Andrea Dell'Amico

- Related to Task #131: Report the R Interpreter environment to reproduce on the Generic Worker nodes added

#2 - May 23, 2015 02:10 PM - Andrea Dell'Amico

- Related to Task #138: Automate the installation of the R suite and its packages added

#3 - May 23, 2015 02:11 PM - Andrea Dell'Amico

- Related to Task #139: Automate the smartgears installation and configuration added

#4 - May 23, 2015 02:13 PM - Andrea Dell'Amico

The distribution must be Ubuntu precise as a R requirement.

Update, from Tom:

The new nodes are up and running. The hostnames go from node43.d4science.org to node78.d4science.org

And this is the nodes distribution on the hypervisor hosts:

```
-dlib14x
node11.d4science.org node35.d4science.org
-dlib15x
node47.d4science.org
-dlib16x
node48.d4science.org
-dlib17x
node34.d4science.org node49.d4science.org
-dlib18x
node55.d4science.org node56.d4science.org
-dlib19x
node12.d4science.org node13.d4science.org
-dlib20x
node38.d4science.org node57.d4science.org node58.d4science.org
-dlib21x
node36.d4science.org node37.d4science.org node59.d4science.org
-dlib22x
node14.d4science.org node15.d4science.org node75.d4science.org
node76.d4science.org
-dlib23x
node16.d4science.org node18.d4science.org node20.d4science.org
node21.d4science.org node23.d4science.org
-dlib24x
node3.d4science.org node4.d4science.org node46.d4science.org
```

```
node73.d4science.org node74.d4science.org
-dlib25x
node50.d4science.org node52.d4science.org node53.d4science.org
node54.d4science.org node60.d4science.org node61.d4science.org
node62.d4science.org node63.d4science.org node77.d4science.org
node78.d4science.org
-dlib26x
node27.d4science.org node28.d4science.org node29.d4science.org
node30.d4science.org node51.d4science.org node64.d4science.org
node65.d4science.org node66.d4science.org node67.d4science.org
node68.d4science.org
-dlib27x
node31.d4science.org node32.d4science.org node33.d4science.org
node43.d4science.org node44.d4science.org node45.d4science.org
node69.d4science.org node70.d4science.org node71.d4science.org
node72.d4science.org
```

#5 - May 23, 2015 03:02 PM - Andrea Dell'Amico

- Assignee changed from Tommaso Piccioli to Andrea Dell'Amico

node43.d4science.org and node44.d4science.org have been deployed with smartgears under the new tomcat package and R. smartgears is configured with the 'dev' scope. I've seen that both register themselves successfully on the d4science dev infrastructure.

If there are no objections, starting on monday morning I'll provision all the new nodes on the production infrastructure.

#6 - May 25, 2015 02:46 PM - Pasquale Pagano

Please update the %Done and let us understand the status of this activity.

#7 - May 25, 2015 03:05 PM - Andrea Dell'Amico

- % Done changed from 0 to 70

The configuration scripts are complete. The R installation has been tested, so I'm launching the configuration of all the new nodes.

#8 - May 25, 2015 03:11 PM - Gianpaolo Coro

I have run precise tests to evaluate the work of this execution environment, both from the Statistical Manager and from R directly. Here are my comments:

- 1 - there are slight changes in the output using the installed version of JAGS, but overall the results are comparable with the previous version and the models converge. Thus, the R environment is OK
- 2 - the other executions on the Statistical Manager using the Worker nodes were successful
- 3 - the GHNs periodically report exceptions, due to some socket timeout and interaction with the Registry. The exceptions are many, one example is :

```
java.lang.IllegalArgumentException: javax.xml.ws.soap.SOAPFaultException
    at org.gcube.informationssystem.publisher.RegistryPublisherImpl.registryUpdate(RegistryPublisherImpl.java:201) ~[registry-publisher-1.2.5-3.7.0.jar:na]
    at org.gcube.informationssystem.publisher.RegistryPublisherImpl.update(RegistryPublisherImpl.java:128) ~[registry-publisher-1.2.5-3.7.0.jar:na]
    at org.gcube.informationssystem.publisher.ScopedPublisherImpl.update(ScopedPublisherImpl.java:54) ~[registry-publisher-1.2.5-3.7.0.jar:na]
    at org.gcube.smartgears.handlers.container.lifecycle.ProfilePublisher.update(ProfilePublisher.java:81) ~[common-smartgears-1.2.2-3.7.0.jar:na]
    at org.gcube.smartgears.handlers.container.lifecycle.ProfileManager.publish(ProfileManager.java:224) [common-smartgears-1.2.2-3.7.0.jar:na]
    at org.gcube.smartgears.handlers.container.lifecycle.ProfileManager.access$300(ProfileManager.java:50) [common-smartgears-1.2.2-3.7.0.jar:na]
    at org.gcube.smartgears.handlers.container.lifecycle.ProfileManager$1.publishAfterChange(ProfileManager.java:122) [common-smartgears-1.2.2-3.7.0.jar:na]
    at sun.reflect.GeneratedMethodAccessor60.invoke(Unknown Source) ~[na:na]
    at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) ~[na:1.7.0_80]
]
    at java.lang.reflect.Method.invoke(Method.java:606) ~[na:1.7.0_80]
    at org.gcube.common.events.impl.Observer.onEventImmediate(Observer.java:99) [common-events-1.0.1-3.7.0.jar:na]
    at org.gcube.common.events.impl.Observer.onEvent(Observer.java:93) [common-events-1.0.1-3.7.0.jar:na]
    at org.gcube.common.events.impl.DefaultHub.notifyObservers(DefaultHub.java:171) [common-events-1.0.1-3.7.0.jar:na]
    at org.gcube.common.events.impl.DefaultHub.fire(DefaultHub.java:93) [common-events-1.0.1-3.7.0.jar:na]
    at org.gcube.smartgears.handlers.container.lifecycle.ProfileManager$2$1.run(ProfileManager.java:275) [common-smartgears-1.2.2-3.7.0.jar:na]
    at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:471) [na:1.7.0_80]
```

```

    at java.util.concurrent.FutureTask.runAndReset(FutureTask.java:304) [na:1.7.0_80]
    at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.access$301(ScheduledThreadPool
Executor.java:178) [na:1.7.0_80]
    at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.run(ScheduledThreadPoolExecuto
r.java:293) [na:1.7.0_80]
    at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) [na:1.7.0_80]
    at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) [na:1.7.0_80]
    at java.lang.Thread.run(Thread.java:745) [na:1.7.0_80]
Caused by: javax.xml.ws.soap.SOAPFaultException: null
    at com.sun.xml.internal.ws.fault.SOAP11Fault.getProtocolException(SOAP11Fault.java:178) ~[na:1.7.0_80]
    at com.sun.xml.internal.ws.fault.SOAPFaultBuilder.createException(SOAPFaultBuilder.java:125) ~[na:1.7.
0_80]
    at com.sun.xml.internal.ws.client.sei.SyncMethodHandler.invoke(SyncMethodHandler.java:108) ~[na:1.7.0_
80]
    at com.sun.xml.internal.ws.client.sei.SyncMethodHandler.invoke(SyncMethodHandler.java:78) ~[na:1.7.0_8
0]
    at com.sun.xml.internal.ws.client.sei.SEIStub.invoke(SEIStub.java:135) ~[na:1.7.0_80]
    at com.sun.proxy.$Proxy38.update(Unknown Source) ~[na:na]
    at org.gcube.informationssystem.publisher.RegistryPublisherImpl.registryUpdate(RegistryPublisherImpl.ja
va:180) ~[registry-publisher-1.2.5-3.7.0.jar:na]
    ... 21 common frames omitted

```

These should be "known" issues, but regard SmartGears and should be solved in the next release.

4 - sometimes, during the download of files from the storage, the GHN on node44 "freezed" for a while. Maybe due to a network problem?

#9 - May 25, 2015 06:17 PM - Andrea Dell'Amico

- Status changed from *In Progress* to *Feedback*

- % Done changed from 70 to 90

All the nodes from node43.d4science.org to node78.d4science.org are now configured with the *production* scope.

The ganglia configuration has been updated too.

#10 - May 26, 2015 06:00 PM - Roberto Cirillo

It's also need to add the production key "d4science.research-infrastructures.eu.gcubekey" on all the new nodes configured for the production scope

#11 - May 26, 2015 11:59 PM - Andrea Dell'Amico

The keys have been installed on all nodes.

#12 - May 29, 2015 12:25 PM - Roberto Cirillo

- Status changed from *Feedback* to *Resolved*

- % Done changed from 90 to 100

#13 - May 29, 2015 12:35 PM - Roberto Cirillo

- Status changed from *Resolved* to *Closed*