

D4Science Infrastructure - Task #10223

Provide machines for FAO course

Nov 06, 2017 05:12 PM - Gianpaolo Coro

Status:	Closed	Start date:	Nov 06, 2017
Priority:	Normal	Due date:	Nov 15, 2017
Assignee:	_InfraScience Systems Engineer	% Done:	100%
Category:	High-Throughput-Computing	Estimated time:	0.00 hour
Target version:	Data Processing		
Infrastructure:	Production		
Description			
Between 21-24 Nov. FAO will do a course using the infrastructure and running CMSY. Since this is a highly demanding process, we should size the computational cluster properly. I would ask to allocate a total of 25 machines for those days. We can try by allocating private-IP machines.			
Related issues:			
Related to D4Science Infrastructure - Task #10358: Dataminer-cluster: Integra...			Closed Nov 16, 2017

History

#1 - Nov 06, 2017 05:36 PM - Andrea Dell'Amico

We can create 11 full sized dataminers at most, before reaching our quota memory limit.

#2 - Nov 07, 2017 07:17 PM - Andrea Dell'Amico

- Status changed from New to In Progress

- % Done changed from 0 to 20

IP and hostnames below, I'll configure them tomorrow. To what cluster do they must be associated? cloud1, with the other GARR instances?

```
192.168.100.3 hostname=dm-192-168-100-3.garr.d4science.org
192.168.100.7 hostname=dm-192-168-100-7.garr.d4science.org
192.168.100.13 hostname=dm-192-168-100-13.garr.d4science.org
192.168.100.14 hostname=dm-192-168-100-14.garr.d4science.org
192.168.100.18 hostname=dm-192-168-100-18.garr.d4science.org
192.168.100.11 hostname=dm-192-168-100-11.garr.d4science.org
192.168.100.12 hostname=dm-192-168-100-12.garr.d4science.org
192.168.100.6 hostname=dm-192-168-100-6.garr.d4science.org
192.168.100.8 hostname=dm-192-168-100-8.garr.d4science.org
192.168.100.9 hostname=dm-192-168-100-9.garr.d4science.org
192.168.100.19 hostname=dm-192-168-100-19.garr.d4science.org
```

#3 - Nov 08, 2017 03:24 PM - Andrea Dell'Amico

The VMs have been provisioned. I'm making them available as single nodes and to the dataminer.garr.d4science.org cluster.

#4 - Nov 08, 2017 03:40 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

- % Done changed from 20 to 100

Can anybody check them?

#5 - Nov 08, 2017 04:05 PM - Gianpaolo Coro

Since the course will be made in the StockAssessment VRE, the machines should be assigned to the cloud1 cluster.

#6 - Nov 08, 2017 04:55 PM - Andrea Dell'Amico

No probl, but should'nt it better thest them before? I'm confident that they are OK, but at least some minimal test would be reassuring.

#7 - Nov 08, 2017 06:42 PM - Pasquale Pagano

I think that it is a good idea to increase the capacities of the cloud1 cluster but please consider that tomorrow there will be the launch of the PAIM VRE that uses that cluster. Since there is a webinar with tens of users registered about PAIM and an intensive promotion of it operated by Trust-IT is

possible that starting from tomorrow there will be an increase in the exploitation of that cluster.

So, if we decide to add those resources to it, it is even better for PAIM but we need to do it now since then we have to retest PAIM.

#8 - Nov 09, 2017 02:49 PM - Andrea Dell'Amico

So, I'll wait for some tests before adding them to the cloud1 backend.

#9 - Nov 10, 2017 12:32 PM - Gianpaolo Coro

- Status changed from Feedback to In Progress

- % Done changed from 100 to 90

Andrea, in order to run our tests we need those machines behind one dedicated cluster. Currently, there is at least another machine in that cluster that is in a prototype scope.

#10 - Nov 10, 2017 02:27 PM - Andrea Dell'Amico

When you confirm that hostname=dm-192-168-100-104.garr.d4science.org can be excluded from the dataminer.garr.d4science.org backend, I'll remove it and you can then use that cluster for your tests.

#11 - Nov 10, 2017 03:02 PM - Gianpaolo Coro

Yes, go on please.

#12 - Nov 10, 2017 03:10 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

Done. The hosts that answer to the dataminer.garr.d4science.org hostname are now

```
- { ip: '192.168.100.3', hostname: 'dm-192-168-100-3.garr.d4science.org' }
- { ip: '192.168.100.7', hostname: 'dm-192-168-100-7.garr.d4science.org' }
- { ip: '192.168.100.13', hostname: 'dm-192-168-100-13.garr.d4science.org' }
- { ip: '192.168.100.14', hostname: 'dm-192-168-100-14.garr.d4science.org' }
- { ip: '192.168.100.18', hostname: 'dm-192-168-100-18.garr.d4science.org' }
- { ip: '192.168.100.11', hostname: 'dm-192-168-100-11.garr.d4science.org' }
- { ip: '192.168.100.12', hostname: 'dm-192-168-100-12.garr.d4science.org' }
- { ip: '192.168.100.6', hostname: 'dm-192-168-100-6.garr.d4science.org' }
- { ip: '192.168.100.8', hostname: 'dm-192-168-100-8.garr.d4science.org' }
- { ip: '192.168.100.9', hostname: 'dm-192-168-100-9.garr.d4science.org' }
- { ip: '192.168.100.19', hostname: 'dm-192-168-100-19.garr.d4science.org' }
```

while dm-192-168-100-104.garr.d4science.org can be reached using its hostname.

#13 - Nov 16, 2017 06:26 PM - Gianpaolo Coro

@lucio.lelii@isti.cnr.it could you run the tests please? The machines are required by tomorrow.

#14 - Nov 16, 2017 07:10 PM - Giancarlo Panichi

- File FirstTest.pdf added

@gianpaolo.coro@isti.cnr.it , I did a fast test on some machines and these are the first results.

#15 - Nov 18, 2017 02:20 PM - Gianpaolo Coro

All the functional tests will be run by Lucio @lucio.lelii@isti.cnr.it on Monday. Meanwhile, for another task I'm running very demanding (CMSY) computations on those machines that use all the cores for a 40 s processing. After about 1000 computations I see that most of the private IP machines are not computing anymore (errors are returned). I cannot see the logs (it seems to me that the logging level is OFF). So I would need help to understand what's happening.

#16 - Nov 19, 2017 07:33 PM - Andrea Dell'Amico

I checked some of them and I didn't find any tomcat/smartgears related exceptions into the log files. The dataminer logs are present into the analysis.log file, accessible from the /gcube-logs/ URI.

Example: <http://dm-192-168-100-14.garr.d4science.org/gcube-logs/> where you can find among the others:

```
[...]
analysis.2017-11-08.0.log      08-Nov-2017  14:16      0
analysis.2017-11-09.0.log      09-Nov-2017  11:58     1564
analysis.2017-11-10.0.log      10-Nov-2017  12:08     7825
```

analysis.2017-11-16.0.log	16-Nov-2017 20:56	466460
analysis.2017-11-17.0.log	17-Nov-2017 19:40	934084
analysis.2017-11-18.0.log	18-Nov-2017 12:16	10488137
analysis.log	18-Nov-2017 15:04	1639399
[...]		

#17 - Nov 19, 2017 09:07 PM - Gianpaolo Coro

The issue is the following and the logs are plenty of them:

```
15:04:46.339 [pool-9-thread-4] DEBUG GenericRScript: Copying /home/gcube/tomcat/webapps/wps/config/./ecocfg/rscr_74258428-a217-4d2c-bc09-c4ec710faa84/SAI_CMSY_FAST/Out_November182017_ID_file.csv.txt to /home/gcube/tomcat/webapps/wps/config/./ecocfg/Out_November182017_ID_file.csv.txt
15:04:46.340 [pool-9-thread-4] ERROR GenericRScript: error in moving file /home/gcube/tomcat/webapps/wps/config/./ecocfg/rscr_74258428-a217-4d2c-bc09-c4ec710faa84/SAI_CMSY_FAST/Out_November182017_ID_file.csv.txt to /home/gcube/tomcat/webapps/wps/config/./ecocfg/Out_November182017_ID_file.csv.txt
org.apache.commons.io.FileExistsException: Destination '/home/gcube/tomcat/webapps/wps/config/./ecocfg/Out_November182017_ID_file.csv.txt' already exists
```

I fear there is a wrong configuration of some path (perhaps the persistence folder of the DM?) because there is an attempt to write the outputs in the ecocfg folder. Another explanation could be that those DataMiners have delays in uploading files to the Workspace and thus concurrent files with the same name are wrongly managed. It needs @lucio.lelii@isti.cnr.it to give it a look to check the configuration and if it gets more complicated we can see this together. I would avoid adding these machines to the cluster before we check them.

So please @lucio.lelii@isti.cnr.it, run the functional tests to see if the DMs work in terms of resources and basic configuration then we investigate this issue.

#18 - Nov 19, 2017 09:38 PM - Andrea Dell'Amico

Apart from the public/private IP address, there is no other difference between those dataminer and the other GARR ones (and between the CNR ones, that have less resources). If there's a problem on them, the same problem is affecting all the dataminer instances.

#19 - Nov 19, 2017 10:06 PM - Gianpaolo Coro

I have just run more than 1000 executions of the CMSY algorithm on dataminer7-p-d4s.d4science.org and I cannot run even one successful execution on those machines. There is something.

#20 - Nov 20, 2017 01:00 PM - Andrea Dell'Amico

As the playbook didn't change and from what I can see the algorithms installer is working, I guess it's better that @lucio.lelii@isti.cnr.it investigates the dataminer behaviour.

#21 - Nov 20, 2017 06:47 PM - Lucio Lelii

@gianpaolo.coro@isti.cnr.it if you are talking about CMSY_2_FAST, I think there is an error in the RScript.

It is trying to write a file with a name created with Month+Day+Year in a common directory, two executions of the same script in the same day in the same server cannot work.

This is the error:

```
8:42:03.957 [pool-9-thread-10] DEBUG GenericRScript: Copying /home/gcube/tomcat/webapps/wps/config/./ecocfg/rscr_5f5a5420-a88c-4193-9856-956ec9d4ca25/SAI_CMSY_FAST/Out_November202017_ID_file.csv.txt to /home/gcube/tomcat/webapps/wps/config/./ecocfg/Out_November202017_ID_file.csv.txt 18:42:03.957 [pool-9-thread-10] ERROR GenericRScript: error in moving file /home/gcube/tomcat/webapps/wps/config/./ecocfg/rscr_5f5a5420-a88c-4193-9856-956ec9d4ca25/SAI_CMSY_FAST/Out_November202017_ID_file.csv.txt to /home/gcube/tomcat/webapps/wps/config/./ecocfg/Out_November202017_ID_file.csv.txt org.apache.commons.io.FileExistsException: Destination '/home/gcube/tomcat/webapps/wps/config/./ecocfg/Out_November202017_ID_file.csv.txt' already exists
```

#22 - Nov 20, 2017 07:08 PM - Gianpaolo Coro

@lucio.lelii@isti.cnr.it that log is from the DataMiner not from the script (see the class name).

The R script writes the file to /ecocfg/rscr_5f5a5420-a88c-4193-9856-956ec9d4ca25/SAI_CMSY_FAST/Out_November202017_ID_file.csv.txt to

Then the DataMiner is autonomously copying it to ecocfg/ and then there is the exception.

On the DM there was an old mechanism that moved a file to a common /persistence/ directory to make it available for uploading to the DataSpace. This mechanism added a UUID and, even when it crashed, it did not influence the execution. I don't think this is involved but the error sounds like there was something similar.

The DM is for some reason copying the file to the cfg directory.

Please, also consider that the same algorithm, with the same code (which is on the WS), works on dataminer7. You can use that as a reference.

#23 - Nov 20, 2017 07:58 PM - Lucio Lelii

I just checked the version of the Dataminer core libraries in both the installation, and they are the same.
I cannot understand how it is happening

#24 - Nov 20, 2017 08:36 PM - Giancarlo Panichi

I do not know if this is important however I noticed that there are different directories in ecocfg:

dm-192-168-100-3

```
./tomcat/webapps/wps/ecocfg/rscr_6aa7a9e6-246d-44ef-93e2-ef0e40552174/SAI_CMSY/  
./tomcat/webapps/wps/ecocfg/rscr_66777071-bccc-4e97-970b-bc866e31285c/SAI_CMSY_FAST/
```

dataminer7-p-d4s

```
./tomcat/webapps/wps/ecocfg/rscr_0235ae5b-0fab-4b5e-9bba-a45c3c47de17/SAI_CMSY_DLM/  
./tomcat/webapps/wps/ecocfg/rscr_69c94c19-26c0-4ce6-bce2-6ca5ee909dbd/SAI_CMSY_FAST/
```

I do not understand if different classes are used

#25 - Nov 21, 2017 02:19 PM - Giancarlo Panichi

@lucio.lelii@isti.cnr.it , I checked the packages, I noticed that the sizes are different, so the builds are different if they have the same filesystem:

dataminer7-p-d4s

```
-rw-r--r-- 1 gcube gcube 81393 Oct 28 04:15 ecological-engine-external-algorithms-1.2.0-4.8.0-132288.jar  
-rw-r--r-- 1 gcube gcube 410909 Oct 28 04:13 ecological-engine-geospatial-extensions-1.5.0-4.8.0-151494.jar  
-rw-r--r-- 1 gcube gcube 248089 Oct 28 00:38 ecological-engine-smart-executor-1.6.0-4.8.0-154627.jar
```

dm-192-168-100-3

```
-rw-r--r-- 1 gcube gcube 81392 Nov 8 04:43 ecological-engine-external-algorithms-1.2.0-4.8.0-132288.jar  
-rw-r--r-- 1 gcube gcube 410908 Nov 8 04:41 ecological-engine-geospatial-extensions-1.5.0-4.8.0-151494.jar  
-rw-r--r-- 1 gcube gcube 248086 Nov 8 00:55 ecological-engine-smart-executor-1.6.0-4.8.0-154627.jar
```

#26 - Nov 22, 2017 01:19 PM - Andrea Dell'Amico

Any progress on this? The deadline was missed already.

#27 - Nov 24, 2017 07:13 PM - Andrea Dell'Amico

- Status changed from Feedback to Closed

- % Done changed from 90 to 100

I'm closing the ticket. Open another one when you feel ready to add them to the cloud1 cluster.

#28 - Nov 24, 2017 09:48 PM - Gianpaolo Coro

There is a general issue on the DataMiner with the management of some concurrent algorithms, highlighted by a recent modification on the SAI. The required modification is two lines of code in the DM, I have communicated it to Lucio. It simply consists of adding a timestamp prefix to the output files of the R scripts.

The real issue on these machines is that the files of the previous computations are not deleted. I'm not able to understand why, also because I cannot access the machines.

I would need to run the following command:

```
for i in `ls -l /proc/*/fd/* 2>/dev/null | grep delete | grep tomcat | awk '{print $9}'`; do du -hL $i | awk '{  
{print $1}' | tr '\n' ' '; ls -l $i | awk '{print $6\" \"$7\" \"$8\" \"$9\" \"$10\" \"$11\" \"$12}'; done
```

#29 - Nov 25, 2017 12:25 PM - Andrea Dell'Amico

A working version of the command:

```
for i in `ls -l /proc/*/fd/* 2>/dev/null | grep delete | grep tomcat | awk '{print $9}'`; do du -hL $i | awk '{  
{print $1}' | tr '\n' ' '; ls -l $i | awk '{print $6 \" \" $7 \" \" $8 \" \" $9 \" \" $10 \" \" $11 \" \" $12}' ; done
```

Empty output on all instances but the following

on dm-192-168-100-11.garr.d4science.org:

```
0 Nov 16 20:11 /proc/23148/fd/917 -> /home/gcube/tomcat/webapps/wps/ecocfg/raster-1465493226242.nc (deleted)
```

on dm-192-168-100-11.garr.d4science.org:

0 Nov 16 19:36 /proc/11111/fd/902 -> /home/gcube/tomcat/webapps/wps/ecocfg/raster-1465493226242.nc (deleted)

on dm-192-168-100-19.garr.d4science.org:

0 Nov 18 01:42 /proc/16711/fd/927 -> /home/gcube/tomcat/webapps/wps/ecocfg/raster-1465493226242.nc (deleted)

A general consideration about the dataminer behaviour. The temporary results of the computations - and the final ones, btw - should be stored outside the webapps directory (this is not only valid for the dataminer, btw). More than that, from the execution logs I've seen on the last days it seems that there isn't any mechanism in place that prevents an algorithms from executing whatever command as gcube user. I don't have to remark how much it's dangerous.

#30 - Nov 28, 2017 04:35 PM - Andrea Dell'Amico

As those VMs were requested to face the FAO course, can we destroy some of them?

#31 - Nov 29, 2017 12:31 PM - Andrea Dell'Amico

After talking with GP we decided to finally add those VMs to the cloud1 cluster. We also destroy some of them: two from the list above, and the one used for the first tests, [#10358](#)
This way we will free some resources in the GARR cloud.

#32 - Nov 29, 2017 12:32 PM - Andrea Dell'Amico

- Related to Task #10358: Dataminer-cluster: Integrate the dm-192-168-100-104.garr.d4science.org instance into the garr cluster added

Files

FirstTest.pdf	31.5 KB	Nov 16, 2017	Giancarlo Panichi
---------------	---------	--------------	-------------------