

## D4Science Infrastructure - Incident #11570

### Algorithms installed in the last weeks are not working

Apr 04, 2018 10:30 PM - Gianpaolo Coro

<b>Status:</b>	Closed	<b>Start date:</b>	Apr 04, 2018
<b>Priority:</b>	Urgent	<b>Due date:</b>	
<b>Assignee:</b>	_InfraScience Systems Engineer	<b>% Done:</b>	100%
<b>Category:</b>	High-Throughput-Computing	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	UnSprintable		
<b>Infrastructure:</b>	Production		
<b>Description</b>			
<p>The algorithms I have installed in the last weeks have suddenly disappeared from the DMs. Is this due to the last upgrade? The error (a null pointer exception when searching for the algorithms) is the one occurring when either the algorithms jar does not exist or has not been installed:</p> <p>Example:</p> <pre>http://dataminer-prototypes.d4science.org/wps/WebProcessingService?request=Execute&amp;service=WPS&amp;Version=1.0.0&amp;gcube-token=&lt;token&gt;&amp;lang=en-US&amp;Identifier=org.gcube.dataanalysis.wps.statisticalmanager.synchserver.mappedclasses.transducerers.LANGUAGE_RECOGNIZER&amp;DataInputs=sentence=North+Korea+has+a+greed+to+send+a+delegation+to+next+month+Winter+Olympics+in+South+Korea%2C+the+first+notable+break+through+to+come+out+of+a+face-to-face+meeting+Tuesday+between+the+neighboring+nations.;</pre>			

### History

#### #1 - Apr 05, 2018 12:06 AM - Pasquale Pagano

- Priority changed from Normal to Urgent

#### #2 - Apr 05, 2018 10:21 AM - Giancarlo Panichi

- Assignee changed from Roberto Cirillo to Massimiliano Assante

Please @massimiliano.assante@isti.cnr.it, the latest version of the SAI:

statistical-algorithms-importer-1.10.0-4.11.0-165574

It seems not to have been installed in production.

You can install this version, thanks.

#### #3 - Apr 05, 2018 10:29 AM - Roberto Cirillo

- Status changed from New to In Progress

The error is the following:

```
2018-04-05 10:10:19,974 [pool-10-thread-5] ERROR AbstractEcologicalEngineMapper: Error execution Algorithm LANGUAGE_RECOGNIZER
java.lang.NullPointerException: null
    at org.gcube.dataanalysis.ecoengine.processing.factories.TransducerersFactory.getTransducerParameters(TransducerersFactory.java:41)
    at org.gcube.dataanalysis.wps.statisticalmanager.synchserver.mapping.AbstractEcologicalEngineMapper.getInputParameters(AbstractEcologicalEngineMapper.java:146)
    at org.gcube.dataanalysis.wps.statisticalmanager.synchserver.mapping.AbstractEcologicalEngineMapper.run(AbstractEcologicalEngineMapper.java:383)
    at org.gcube.dataanalysis.wps.statisticalmanager.synchserver.mappedclasses.transducerers.LANGUAGE_RECOGNIZER.run(LANGUAGE_RECOGNIZER.java:24)
    at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
    at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62)
    at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43)
    at java.lang.reflect.Method.invoke(Method.java:498)
    at org.n52.wps.algorithm.annotation.AnnotationBinding$ExecuteMethodBinding.execute(AnnotationBinding.java:89)
    at org.n52.wps.server.AbstractAnnotatedAlgorithm.run(AbstractAnnotatedAlgorithm.java:54)
    at org.gcube.data.analysis.wps.ExecuteRequest.call(ExecuteRequest.java:608)
```

```

at org.gcube.data.analysis.wps.ExecuteRequest.call(ExecuteRequest.java:67)
at org.gcube.common.authorization.library.AuthorizedTasks$1.call(AuthorizedTasks.java:41)
at java.util.concurrent.FutureTask.run(FutureTask.java:266)
at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1149)
at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:624)
at java.lang.Thread.run(Thread.java:748)

```

The algorithm entry is the following:

```

| LANGUAGE_RECOGNIZER | Gianpaolo Coro | NATURAL_LANGUAGE_PROCESSING | DataMinerPoolManager | <notextile>./add
Algorithm.sh LANGUAGE_RECOGNIZER NATURAL_LANGUAGE_PROCESSING org.gcube.dataanalysis.executor.rscripts.LANGUAGE
RECOGNIZER /d4science.research-infrastructures.eu/gCubeApps/RPrototypingLab transducerers N http://data.d4scie
nce.org/aGovM2Q5TThJaUtyak9ab0dFT2J6V2RveGVMazlWN0NHbWJQNStIS0N6Yz0 "This algorithm identifies the language of
an input sentence." </notextile> | none | Mon Feb 12 11:01:44 UTC 2018 |

```

and the jar seems to be there:

```

gcube@dataminer2-proto:~$ ls -als wps_algorithms/algorithms/proto/software/LANGUAGE_RECOGNIZER*
4 -rw-rw-r-- 1 gcube gcube 2761 Feb 12 12:02 wps_algorithms/algorithms/proto/software/LANGUAGE_RECOGNIZER.jar
4 -rw-rw-r-- 1 gcube gcube 1782 Feb 12 12:02 wps_algorithms/algorithms/proto/software/LANGUAGE_RECOGNIZER_inte
rface.jar

```

Everything seems to be correct by my side. What is wrong?

#### #4 - Apr 05, 2018 10:48 AM - Gianpaolo Coro

Then the algorithm is not being installed.

#### #5 - Apr 05, 2018 10:55 AM - Giancarlo Panichi

@roberto.cirillo@isti.cnr.it , it may depend on the installation of the SAI which needs to be updated.

#### #6 - Apr 05, 2018 11:07 AM - Roberto Cirillo

- Status changed from In Progress to Feedback

- Assignee changed from Massimiliano Assante to Gianpaolo Coro

Gianpaolo Coro wrote:

Then the algorithm is not being installed.

Please, could you explain better? Is it the right algorithm? As you can see, the algorithm is present and it has been installed the 12th Feb. When have you performed the last algorithm update?

#### #7 - Apr 05, 2018 11:16 AM - Massimiliano Assante

- Status changed from Feedback to In Progress

- Assignee changed from Gianpaolo Coro to Massimiliano Assante

- % Done changed from 0 to 20

The SAI deployed was the one build on the night of March 23rd, apparently after that day it was updated again (March 29th) without advising me:

The related ticket release #11490 <https://support.d4science.org/issues/11490> hasn't been updated since 15 days (the day I deployed the version on preprod). Next time make sure you update it if the preprod version changes.

The SW Test ticket <https://support.d4science.org/issues/11491> had everyone but me on copy (assuming that by reading it one could understand that the war was updated)..

As you can see for me it was impossible to know.

Anyhow I'm now going to update the Latest SAI on production.

#### #8 - Apr 05, 2018 11:18 AM - Massimiliano Assante

typo before, I meant I'm now going to update Latest SAI on production.

#### #9 - Apr 05, 2018 11:21 AM - Massimiliano Assante

- % Done changed from 20 to 50

statistical-algorithms-importer-1.10.0-4.11.0-165574 upgrade on production is in progress

**#10 - Apr 05, 2018 11:38 AM - Massimiliano Assante**

- Assignee changed from Massimiliano Assante to Giancarlo Panichi
- % Done changed from 50 to 70

statistical-algorithms-importer-1.10.0-4.11.0-165574 upgrade on production is completed

**#11 - Apr 05, 2018 11:42 AM - Giancarlo Panichi**

- Assignee changed from Giancarlo Panichi to Gianpaolo Coro

Perfect @massimiliano.assante@isti.cnr.it , now the SAI works properly.

Please @gianpaolo.coro@isti.cnr.it , can you republish your algorithm so we can check that everything works?  
Thanks.

**#12 - Apr 05, 2018 11:44 AM - Gianpaolo Coro**

- Assignee changed from Gianpaolo Coro to Giancarlo Panichi

Please don't write about SAI in this ticket. It is about algorithms there were installed weeks ago and have never been updated recently, no touched, and now simply have disappeared. This is the case of many algorithms, I have just reported one example. If a new DM installation was done, the algorithms are erased and installed from scratch. Is this the case? Perhaps something went wrong in the process?

**#13 - Apr 05, 2018 11:46 AM - Gianpaolo Coro**

Obviously I cannot re-publish all the algorithms, because I'm not the author of many of them. The example I have reported should work as a guidance to figure out the issue.

**#14 - Apr 05, 2018 11:59 AM - Andrea Dell'Amico**

Gianpaolo Coro wrote:

If a new DM installation was done, the algorithms are erased and installed from scratch. Is this the case? Perhaps something went wrong in the process?

No.

**#15 - Apr 05, 2018 12:09 PM - Giancarlo Panichi**

- Assignee changed from Giancarlo Panichi to Gianpaolo Coro

@gianpaolo.coro@isti.cnr.it , I can only say that this algorithm LANGUAGE\_RECOGNIZER has not been installed correctly because information on userperspective.properties and transducerers.properties files are missing.

The only thing you can do is republish it.

The other installations do not seem to have problems, however if you try the other algorithms you can tell us which ones do not work and we can investigate.

**#16 - Apr 05, 2018 12:17 PM - Giancarlo Panichi**

@gianpaolo.coro@isti.cnr.it ,however there is something true in what you say given that now there are 93 algorithms in the category Others seems a bit 'much.

Maybe the userperspective.properties and transducerers.properties files were changed in the upgrade?

**#17 - Apr 05, 2018 12:24 PM - Andrea Dell'Amico**

Giancarlo Panichi wrote:

@gianpaolo.coro@isti.cnr.it ,however there is something true in what you say given that now there are 93 algorithms in the category Others seems a bit 'much.

Maybe the userperspective.properties and transducerers.properties files were changed in the upgrade?

Obviously yes. There's a list of files that overwrite the dataminer defaults and that they are installed after every upgrade. The files are:

```
- algorithms.properties
- clusterers.properties
- dynamictransducerers.properties
- evaluators.properties
- generators.properties
- modelers.properties
- models.properties
- nodealgorithms.properties
```

```
- transducerers.properties
- userperspective.properties
```

If they are meant to change when some algorithms are installed, they must be managed in a different way. And moved outside the tomcat/webapps/wps directory.

**#18 - Apr 05, 2018 12:33 PM - Giancarlo Panichi**

@andrea.dellamico@isti.cnr.it , when you update the DMs, you also install the algorithms and then update those files.

If all went well, the files should contain the correct information.

In this case of LANGUAGE\_RECOGNIZER the two files above do not contain the information, so something was not successful in the update.

Are all the algorithms reinstalled after a service update? Could this be due to the number of algorithms that are now so many? Maybe some conflict arises?

Perhaps in the event of an update it may be appropriate to choose a different policy.

**#19 - Apr 05, 2018 01:07 PM - Giancarlo Panichi**

@lucio.lelii@isti.cnr.it , it may be necessary to move these files externally as @andrea.dellamico@isti.cnr.it suggests, and to evaluate if perhaps it would be better if they were a resource on the IS.

**#20 - Apr 05, 2018 01:41 PM - Andrea Dell'Amico**

Giancarlo Panichi wrote:

@andrea.dellamico@isti.cnr.it , when you update the DMs, you also install the algorithms and then update those files.

No algorithm is installed during an update, the algorithms file outside the application.

Only the files listed above are provisioned and they overwrite the ones provided by the dataminer war file. But if they are changed when algorithms are installed, this isn't the correct procedure.

If all went well, the files should contain the correct information.

In this case of LANGUAGE\_RECOGNIZER the two files above do not contain the information, so something was not successful in the update.

Are all the algorithms reinstalled after a service update? Could this be due to the number of algorithms that are now so many? Maybe some conflict arises?

Perhaps in the event of an update it may be appropriate to choose a different policy.

Nobody ever told me that an upgrade must trigger a reinstall of all the algorithms. It's an easy operation btw, everybody can trigger it just removing or emptying the wps\_algorithms\_install\_log/already\_installed\_algorithms.txt file.

Let me know if I have to do it.

**#21 - Apr 05, 2018 01:45 PM - Gianpaolo Coro**

These are the FACTS we should consider (I have done specific verification on dataminer3-proto):

1 - several algorithms were installed in March and they have disappeared from the configuration files. The only way to do this is to completely overwrite the files, which I have never done, obviously and has been done by some script. This script could be the installation, the update or something else I don't know. BUT, these files have been overwritten.

2 - I don't own all these algorithms, I don't have the complete list of the ones missing, thus I cannot install them through SAI. Several of them are also in other VREs than RProto, thus reinstallation through SAI is not feasible!

3 - The algorithms are correctly installed if I run the installation instruction contained in the DMPool Manager string. Thus no issue on the installer

Thus, after this rational analysis the most probable issues are in the fact that either the installation of all the algorithms was not run or the installation cron is not working. Thus, I would force the run of the cron job to exclude the second.

Generally, If I have correctly understood the new version produced by @lucio.lelii@isti.cnr.it , this starts from an empty DataMiner and requires all the algorithms to be installed from scratch by triggering the cron job.

**#22 - Apr 05, 2018 01:48 PM - Andrea Dell'Amico**

Giancarlo Panichi wrote:

@lucio.lelii@isti.cnr.it , it may be necessary to move these files externally, and to evaluate if perhaps it would be better if they were a resource on the IS.

I guess that an IS resource would be the best option.

Other than that, everything that is not part of the original dataminer war file (only exception the basic configuration that never changes after the installation, like wps\_config.xml) should live elsewhere.

There is a lot of stuff that happens inside tomcat/webapps/wps right now, and shouldn't:

- algorithms execution temporary files and results
- property files
- gcube keys (apparently the ones already present inside tomcat/lib are not enough)
- The gebco\_08.nc (I managed to move it outside and use a symbolic link, otherwise every dataminer upgrade would download it again, and it's almost 2GB).
- Maybe other that I don't know about.

**#23 - Apr 05, 2018 01:54 PM - Andrea Dell'Amico**

Gianpaolo Coro wrote:

These are the FACTS we should consider (I have done specific verification on dataminer3-proto):

1 - several algorithms were installed in March and they have disappeared from the configuration files. The only way to do this is to completely overwrite the files, which I have never done, obviously and has been done by some script. This script could be the installation, the update or something else I don't know. BUT, these files have been overwritten.

2 - I don't own all these algorithms, I don't have the complete list of the ones missing, thus I cannot install them through SAI. Several of them are also in other VREs than RProto, thus reinstallation through SAI is not feasible!

3 - The algorithms are correctly installed if I run the installation instruction contained in the DMPool Manager string. Thus no issue on the installer

Thus, after this rational analysis the most probable issues are in the fact that either the installation of all the algorithms was not run or the installation cron is not working. Thus, I would force the run of the cron job to exclude the second.

@gianpaolo.coro@isti.cnr.it read above: the property files **are** overwritten at each dataminer upgrade, because that's what I was told to do. The cron job has nothing to do with that, it executes the addAlgorithm script and we know that it works correctly (but for the [ ] case).

Generally, If I have correctly understood the new version produced by @lucio.lelii@isti.cnr.it , this starts from an empty DataMiner and requires all the algorithms to be installed from scratch by triggering the cron job.

The part where all the algorithms to be installed from scratch never arrived to me. Do you confirm that running the script executed from the cron job should fix?

**#24 - Apr 05, 2018 02:02 PM - Gianpaolo Coro**

@lucio.lelii@isti.cnr.it do you confirm that the algorithms installation should be triggered just after each re-installation?

@andrea.dellamico@isti.cnr.it please, run the cron job and let's see if everything works, because this issue is blocking several tools that build on these functionalities.

Sorry, perhaps I lost some answer, since this incident ticket is wrongly collecting also SAI requests, suggestions for DM best practices, and lifestyle advice. Let's keep all of this in separate tickets, please.

**#25 - Apr 05, 2018 02:13 PM - Andrea Dell'Amico**

- Assignee changed from Gianpaolo Coro to \_InfraScience Systems Engineer

OK, I'm forcing the removal of /home/gcube/wps\_algorithms\_install\_log/already\_installed\_algorithms.txt

**#26 - Apr 05, 2018 02:22 PM - Andrea Dell'Amico**

Removed the file and forced a run of /usr/local/bin/algorithms-updater on all the dev and preprod dataminers. The job is currently running on the production ones four at a time, it should complete in a couple of minutes. Then you'll have to wait for the algorithms installation procedure to complete.

**#27 - Apr 05, 2018 02:22 PM - Andrea Dell'Amico**

- Status changed from In Progress to Feedback

- % Done changed from 70 to 100

**#28 - Apr 05, 2018 02:24 PM - Gianpaolo Coro**

- Status changed from Feedback to In Progress

- % Done changed from 100 to 70

I don't know how this "/home/gcube/wps\_algorithms\_install\_log/already\_installed\_algorithms.txt" works, what it contains, and when it is generated.

Let's wait for an answer by @lucio.lelii@isti.cnr.it

**#29 - Apr 05, 2018 02:26 PM - Andrea Dell'Amico**

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

- % Done changed from 70 to 100

I just talked with Lucio, FYI.

It's the file that lists the already installed algorithms, as the name suggests. Removing it, the addAlgorithm should be run on every listed algorithm (I have the doubt that it exits without doing anything if the svn update does not show any update. Can you redeploy one of the algorithms, just to update the subversion repository?).

**#30 - Apr 05, 2018 02:34 PM - Andrea Dell'Amico**

I confirm that svn must be updated to trigger the script execution. If it can't be done, I'm going to change the script so that it will run unconditionally when the /home/gcube/wps\_algorithms\_install\_log/already\_installed\_algorithms.tx is absent.

**#31 - Apr 05, 2018 02:58 PM - Gianpaolo Coro**

- Assignee changed from *\_InfraScience Systems Engineer* to *Lucio Lelii*

- % Done changed from 100 to 70

Ok, I'm updating the SVN.

**#32 - Apr 05, 2018 03:01 PM - Gianpaolo Coro**

I have changed one creation time in the algorithms list and committed both in prod and in proto. This should trigger the installation if I have correctly understood.

**#33 - Apr 05, 2018 03:08 PM - Andrea Dell'Amico**

You missed the 15:00 cron job run by a minute, so I triggered a run of the script.

**#34 - Apr 05, 2018 03:09 PM - Andrea Dell'Amico**

- Assignee changed from *Lucio Lelii* to *\_InfraScience Systems Engineer*

- % Done changed from 70 to 100

I'm going to open a new ticket to list all the problems with the current implementation.

**#35 - Apr 05, 2018 04:13 PM - Gianpaolo Coro**

- Status changed from *Feedback* to *Resolved*

The algorithms are working again. Thank you.

**#36 - Apr 05, 2018 04:23 PM - Andrea Dell'Amico**

- Status changed from *Resolved* to *Closed*