

D4Science Infrastructure - Support #11605

Set DataMiner queue length to 1

Apr 10, 2018 05:42 PM - Gianpaolo Coro

Status:	Closed	Start date:	Nov 02, 2018
Priority:	Urgent	Due date:	
Assignee:	Leonardo Candela	% Done:	100%
Category:	High-Throughput-Computing	Estimated time:	0.00 hour
Infrastructure:	Production		
Description In the templates/setup.cfg file on the DataMiner, the maxcomputations parameter should be changed to 1 instead of 4. This parameter should be changed in the component too, possibly. However, this modification is urgent to users currently using DataMiner as a cloud computing platform, thus action via provisioning is required.			
Subtasks:			
Task # 12614: Need server with joblength 1 to use all 16 cores and a cue system per ser...			Closed
Task # 12806: Testing the new settings of the RAKIP infrastructure			Closed

History

#1 - Apr 10, 2018 07:40 PM - Andrea Dell'Amico

Where is it?

```
root@dataminer1-d-d4s:/home/gcube/tomcat/webapps/wps# find . -name setup.cfg
root@dataminer1-d-d4s:/home/gcube/tomcat/webapps/wps#
```

#2 - Apr 10, 2018 07:45 PM - Gianpaolo Coro

@lucio.lalii@isti.cnr.it could you help us, please?

#3 - Apr 11, 2018 10:29 AM - Lucio Lalii

It is in src/main/resources inside the dataminer jar, It cannot be changed in this way

#4 - Apr 11, 2018 11:56 AM - Lucio Lalii

I'm building, via etics, the dataminer jar with the modification needed. When it will be ready we need to replace the jar in every dataminer.

#5 - Apr 11, 2018 04:20 PM - Andrea Dell'Amico

Please let me know when it's ready and post the maven coordinates, so that I can put in place an ad hoc playbook.

#6 - Apr 12, 2018 12:39 PM - Andrea Dell'Amico

- Status changed from New to In Progress

We have the artifact. I'm going to install it on the dev dataminers first.

#7 - Apr 12, 2018 01:01 PM - Andrea Dell'Amico

Can anybody check dataminer1-d-d4s.d4science.org? If I do not hear anything back in an hour, I'll start updating the production dataminers.

#8 - Apr 12, 2018 02:19 PM - Andrea Dell'Amico

- % Done changed from 0 to 30

The services in dev restarted correctly, I'm starting rolling out in production.

#9 - Apr 12, 2018 02:30 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

- % Done changed from 30 to 100

Done.

#10 - Apr 12, 2018 04:32 PM - Gianpaolo Coro

As I have discussed with Andrea, this patch should be rolled-back since there are algorithms that invoke themselves, possibly on the same machine. In the next weeks, Andrea will deploy new DataMiner machines as Generic Worker machines behind the Generic Worker proxy, and this should solve the issue.

#11 - Apr 12, 2018 04:41 PM - Andrea Dell'Amico

The rollback is running.

Gianpaolo Coro wrote:

As I have discussed with Andrea, this patch should be rolled-back since there are algorithms that invoke themselves, possibly on the same machine. In the next weeks, Andrea will deploy new DataMiner machines as Generic Worker machines behind the Generic Worker proxy, and this should solve the issue.

Would you open a dedicated ticket for the new generic workers? They substitute the old ones, correct?
And when deployed, the queue lenght parameter will be fixed at 1 on both the regular dataminer and the generic workers?

#12 - Apr 12, 2018 04:42 PM - Andrea Dell'Amico

(rollback completed)

#13 - Apr 12, 2018 05:28 PM - Gianpaolo Coro

Thank you. Given the way people are going to use the services, I think we will need to have the standard DataMiners running with the default configuration and the Generic Workers running with 1.

#14 - Apr 12, 2018 05:30 PM - Andrea Dell'Amico

Gianpaolo Coro wrote:

Thank you. Given the way people are going to use the services, I think we will need to have the standard DataMiners running with the default configuration and the Generic Workers running with 1.

So that value must be converted in a property, configurable at provisioning time without swapping jar files.

#15 - Apr 12, 2018 05:53 PM - Gianpaolo Coro

- Status changed from Feedback to Closed

#16 - Sep 01, 2018 12:35 PM - Lars Valentin

- Tracker changed from Task to Support

- Assignee changed from _InfraScience Systems Engineer to Leonardo Candela

- Priority changed from High to Normal

Hello Leonardo,

I hope you don't mind that I reuse this specific ticket instead of opening a new one with the same context.

Unfortunately, I did not use the chance in the webmeeting yesterday, since, I would like to know the current status of the queue length per damaminer server and the future possibilities.

My understanding in the past was that each dataminer has 16 cores and 4 slots to accept jobs - what means that each job is allowed to use max. 4 cores. Am I right at this point?

That is totally fine if we run workflows (models) which are not optimized for multicore use and therefore only use one core.

But what about workflows/models which are able to use all cores provided by the system? Usually, they should be programmed to leave one core for system processes, which would mean 3 cores are available to be used if one job can only assign 4 cores max.

I would like to ask if there is the possibility to execute workflows/models via REST in a way that the dataminer knows it should switch and run this job unique on the 16 cores dataminer to allow to incorporate all 16 cores for one job exclusively.

That would be 15 cores available to run the model instead of 3, what might roughly lead to 1/5 of the previous running time.

Thank you in advance!

Lars

PS. Do you may have dataminer with more cores to be addressed in the future for such cases?

#17 - Sep 06, 2018 12:45 PM - Leonardo Candela

Hi @lars.valentin@bfr.bund.de I would suggest to open specific tickets to discuss any need you / your use cases might have on aspects like that.

During our last PMB meeting (see slides here <https://goo.gl/ZFYGJq>) I tried to explain that:

- right now there are two "clusters" configured to provide data miner facilities (proto and prod);
- these clusters are not for exclusive use of AGINFRA+ cases;
- we can have one cluster per VRE;

The business logic the service use for allocating the tasks on machines and consuming the available cores can be better explained by @gianpaolo.coro@isti.cnr.it ... in particular regarding cores (a) it depends also on how the algorithm has been implemented and (b) it cannot be changed at algorithm invocation time.

Last but not least, if you have particular settings / behaviors you have to satisfy we can try to configure a specific cluster yet this has a cost we have to carefully evaluate.

These are overall comments stemming from me, my suggestion is to try to be specific and report any specific need you have (e.g. enact process x that needs 100 cores) and we will do our best to satisfy it with the available resources and technologies.

#18 - Nov 02, 2018 04:25 PM - Lars Valentin

- *Due date set to Nov 02, 2018*

- *Start date changed from Oct 06, 2018 to Nov 02, 2018*

due to changes in a related task