

D4Science Infrastructure - Task #12614

Support # 11605 (Closed): Set DataMiner queue length to 1

Need server with joblength 1 to use all 16 cores and a cue system per server.

Oct 06, 2018 01:57 PM - Lars Valentin

Status:	Closed	Start date:	Nov 02, 2018
Priority:	Urgent	Due date:	
Assignee:	Lars Valentin	% Done:	100%
Category:	High-Throughput-Computing	Estimated time:	0.00 hour
Target version:	No		
Infrastructure:	Development, Pre-Production, Production		
Description			
as explained here https://support.d4science.org/issues/11605#note-16 I follow the wish from https://support.d4science.org/issues/11605#note-17 and create a new ticket:			
We need server which are able to use all cores of a server at once for a single job (calculating models). Right now the core max is 16 I think (in the future we would like to have more ~100+). The server should also have a cue system in order to accept many jobs, but run only one after the other. Ideal for us would be if all server of a cluster behave like that, in order to avoid the need to adress a single server inside the cluster.			
Subtasks:			
Task # 12806: Testing the new settings of the RAKIP infrastructure			Closed

History

#1 - Oct 08, 2018 04:42 PM - Leonardo Candela

Apart from the technicalities behind this request, there is an organisational issue.

It was clarified many times that there is no DM cluster dedicated to AGINFRA+ use cases, the cluster is shared with other communities. If we change the behavior of such a shared cluster, the change will affect all the other communities. A machine with ~100+ cores is not in the current portfolio of AGINFRA / D4Science ;)

If this request is for the creation of a new cluster please clarify the "scope", i.e. how many processes needs this setting. The creation of a cluster has a cost in terms of human time and machines to allocate. It should be discussed at PMB.

You last comment is quite puzzling to me ... "Ideal for us would be if all server of a cluster behave like that, in order to avoid the need to adress a single server inside the cluster." ... this is definitely a non recommended practice ... the cluster might evolve, the developer of an algorithm should not be exposed to the internals of DataMiner and its management of machines.

#2 - Oct 08, 2018 05:45 PM - Matthias Filter

I agree to discuss this during the next PBM.

I will add this ticket to the WP6 discussion slot.

#4 - Oct 08, 2018 07:15 PM - Lars Valentin

Good idea to discuss that during PBM. I would like to add a few thoughts for later discussion.

I understand that there are other users on the system, which would be influenced by a change. Therefore, it would be interesting what is the advantage on a joblength = 4 for them and whether they would be happy with a joblength=1.

I don't understand why the last comment is puzzling for you. 'to adress a single server inside the cluster' is what I want to avoid, here we are exactly on the same page. But in the past CNR mentioned in a ticket that this is possible via REST and a workaround. All I ask for is a solution for us to make a meaningful use of your infrastructure for our use case (since an HPC is not available). Maybe you can provide some ideas until the meeting?

Just to clarify the scope: Allowing a job(model) to use all cores (the more the better) of a server, without beeing disturbed by another job via cueing system (as it is already implemented).

In my understanding the solution to that could be a joblength=1.

#5 - Oct 30, 2018 10:18 AM - Leonardo Candela

- Tracker changed from Support to Task

- Assignee changed from Leonardo Candela to Roberto Cirillo

According to the discussion we had at the meeting, I kindly ask @roberto.cirillo@isti.cnr.it to change the configuration of the DM cluster serving the RAKIP VRE.

In essence we should configure DM to use directly the "workers" cluster ... this will lead to the "queue 1" behavior.

#6 - Oct 31, 2018 02:43 PM - Roberto Cirillo

- Status changed from New to In Progress

#7 - Oct 31, 2018 02:50 PM - Roberto Cirillo

- Status changed from In Progress to Feedback

- Assignee changed from Roberto Cirillo to Leonardo Candela

- % Done changed from 0 to 100

The dm cluster with queue 1 has been configured on RAKIP_Portal VRE. Please @leonardo.candela@isti.cnr.it could you check if it works as expected?

#8 - Oct 31, 2018 03:37 PM - Leonardo Candela

- Assignee changed from Leonardo Candela to Lars Valentin

I would like to pass the task to @lars.valentin@bfr.bund.de ... the community is not in the position to challenge the cluster and check whether this configuration and behaviour of the DM cluster is suitable for the exploitation scenario.

One option is to perform this activity as follows:

- close this ticket since the cluster is configured as expected;
- open a new ticket declaring the testing / assessment plan (e.g. algorithms to be integrated and tested, load, concurrency, number of users served in parallel) and use the new ticket to report on the effectiveness of the proposed solution.

#9 - Nov 27, 2018 12:30 PM - Leonardo Candela

- Status changed from Feedback to Closed

I'm going to close the ticket. According to [#12806](#) the RAKIP VRE cluster is supporting the expected case.