

D4Science Infrastructure - VM Creation #7927

More DataMiner machines required in the GenericWorker cluster

Mar 31, 2017 04:43 PM - Gianpaolo Coro

| | | | |
|---|---|---------------------------------------|--------------|
| Status: | Closed | Start date: | Mar 31, 2017 |
| Priority: | Normal | Due date: | Apr 07, 2017 |
| Assignee: | _InfraScience Systems Engineer | % Done: | 100% |
| Target version: | Resource Provisioning | | |
| Infrastructure: | Production | Applications list and version: | DataMiner |
| RAM: | 8 | Applications context path: | /wps |
| CPU: | 2 | External storage data disk: | |
| DiskSpace: | 10 | Smartgears auth connector: | |
| Smartgears: | Yes | data transfer enabled: | No |
| Smartgears Version: | Latest | PostgreSQL database: | No |
| Application Instructions: | https://wiki.gcube-system.org/gcube/DataMiner_Installation | Number of VMs: | |
| Smartgears scopes: | Possibly all the production scopes | | |
| Description | | | |
| 10 DataMiner machines are required in the GenericWorker cluster to provision BiOnym as-a-Service for FAO. These machines can be asked to the Garr infrastructure. | | | |

History

#1 - Mar 31, 2017 08:13 PM - Andrea Dell'Amico

- Status changed from New to In Progress
- % Done changed from 0 to 30

Provisioning is running. The dataminers will be available on all the scopes.

#2 - Mar 31, 2017 08:54 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback
- % Done changed from 30 to 100

The VMs are active. hostnames are:

```
ip-90-147-167-173.ct1.garrservices.it
ip-90-147-167-175.ct1.garrservices.it
ip-90-147-167-176.ct1.garrservices.it
ip-90-147-167-177.ct1.garrservices.it
ip-90-147-167-178.ct1.garrservices.it
ip-90-147-167-179.ct1.garrservices.it
ip-90-147-167-180.ct1.garrservices.it
ip-90-147-167-181.ct1.garrservices.it
ip-90-147-167-182.ct1.garrservices.it
ip-90-147-167-183.ct1.garrservices.it
```

#3 - Mar 31, 2017 11:13 PM - Tommaso Piccioli

New nodes added to the iptable rules on workspace-repository.d4science.org
(and the new workspace-repository-prod.d4science.org for the future)

#4 - Apr 07, 2017 03:37 PM - Gianpaolo Coro

- Status changed from Feedback to In Progress

The machines should be added to the genericworkers cluster in order to be properly tested.

#5 - Apr 10, 2017 11:15 AM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

Gianpaolo Coro wrote:

The machines should be added to the genericworkers cluster in order to be properly tested.

Done.

#6 - Apr 12, 2017 03:48 PM - Gianpaolo Coro

The machines work very well and I was able to use them to run 3500 taxonomic searches in a very large repository. I'm just testing the dataminer1-2-genericworker machines that seem to have issues. For sure they are downsized, due also to the scheduled updater of R packages, but it seems they have other issues too with Java libraries. Is it possible to increase their memory configuration and reinstall them?

#7 - Apr 12, 2017 08:03 PM - Andrea Dell'Amico

Gianpaolo Coro wrote:

I'm just testing the dataminer1-2-genericworker machines that seem to have issues. For sure they are downsized, due also to the scheduled updater of R packages, but it seems they have other issues too with Java libraries. Is it possible to increase their memory configuration and reinstall them?

I don't see why a reinstall should be useful. Their resources can be increased, but the distribution and the configurations are up to date with all the other dataminer installations. I don't know if there's some outdated algorithm, but it's not something that a reinstall can solve.

For the time being, if the GARR dataminers are sufficient, the local generic workers can be removed from the haproxy pool and maybe dismissed.

#8 - Apr 19, 2017 07:38 PM - Andrea Dell'Amico

@gianpaolo.coro@isti.cnr.it any opinion on my proposal?

#9 - Apr 19, 2017 11:26 PM - Gianpaolo Coro

Yes, I have run some tests and the Garr resources are sufficient.

#10 - Apr 20, 2017 11:58 AM - Andrea Dell'Amico

I just removed dataminer[1:2]-genericworkers.d4science.org from the genericworkers backend. Now we have two free dataminer VMs. We can add them to the genericworkers_proto backend or we could destroy them and add more resources to dataminer3-genericworkers.d4science.org. @gianpaolo.coro@isti.cnr.it, your call.

#11 - Apr 20, 2017 12:09 PM - Gianpaolo Coro

Please, add resources to dataminer3-genericworkers.d4science.org

#12 - Apr 20, 2017 12:13 PM - Andrea Dell'Amico

Gianpaolo Coro wrote:

Please, add resources to dataminer3-genericworkers.d4science.org

Is it a problem if I increase the dataminer1-genericworkers.d4science.org's resources and put it into the genericworkers_proto backend? It lives in an hypervisor with more free memory than the one that hosts dataminer3-genericworkers.d4science.org

#13 - Apr 20, 2017 03:26 PM - Andrea Dell'Amico

- Status changed from Feedback to Closed

Nevermind. dataminer3-genericworkers.d4science.org has now 8 GB of RAM and 4 CPUs.

I'm going to shutdown dataminer{1,2}-genericworkers.d4science.org