

D4Science Infrastructure - Task #169

Verification of the smartgears nodes list

May 27, 2015 04:08 PM - Andrea Dell'Amico

Status:	Closed	Start date:	May 27, 2015
Priority:	Normal	Due date:	
Assignee:	Roberto Cirillo	% Done:	100%
Category:	System Application	Estimated time:	0.00 hour
Target version:	System Configuration		
Infrastructure:	Production		
Description Ganglia is showing more smartgears workers than expected: [[[http://monitoring.research-infrastructures.eu/ganglia/?c=D4science%20Smartgears%20cluster&m=load_one&r=day&s=by%20name&hc=4&mc=2]]]. Gianpaolo says that we are listing as worker nodes hosts that aren't workers even if they are part of the smartgears cluster. I think that we need a verification of the ghn hosts list that lives in the ansible-playbooks git repo, here: d4science-ghn-cluster/inventory/hosts.production. Note that the new smartgears nodes are listed in a separate file: d4science-ghn-cluster/inventory/hosts.smartgears. The two lists will be merged when all the hosts will be completely provisioned via ansible.			

History

#1 - Jul 30, 2015 12:24 PM - Roberto Cirillo

- Status changed from New to In Progress

The list of new workers defined in "host.smartgears" file is right
But the list of old workers defined in "host.production" is not right:

I see in ganglia all the nodes of the group "ghn_smartgears_prod_old" but not all the nodes presents in this list are workers. For example "tabulardata.d4science.org" is a node smartgears based but this is not a worker.
I see the group "ghn_workers_prod" empty in "hosts.production". Why? Maybe all the workers have to be defined in this group.

In addition, Ganglia is showing workers that are not present in any lists:
node13.p.d4science.research-infrastructures.eu
node31.p.d4science.research-infrastructures.eu
node51.p.d4science.research-infrastructures.eu

#2 - Jul 30, 2015 12:51 PM - Andrea Dell'Amico

Roberto Cirillo wrote:

The list of new workers defined in "host.smartgears" file is right
But the list of old workers defined in "host.production" is not right:

I see in ganglia all the nodes of the group "ghn_smartgears_prod_old" but not all the nodes presents in this list are workers. For example "tabulardata.d4science.org" is a node smartgears based but this is not a worker.

Ok.

I see the group "ghn_workers_prod" empty in "hosts.production". Why? Maybe all the workers have to be defined in this group.

That group is used inside host.smartgears and contains the ansible provisioned hosts. The other ones are under ghn_smartgears_prod_old. When will be ready to also provision the old workers, they all will be gathered under ghn_smartgears_prod and the inventory file host.smartgears will disappear

In addition, Ganglia is showing workers that are not present in any lists:
node13.p.d4science.research-infrastructures.eu
node31.p.d4science.research-infrastructures.eu
node51.p.d4science.research-infrastructures.eu

Ah. I'll remove them too.

#3 - Jul 30, 2015 01:12 PM - Andrea Dell'Amico

- % *Done* changed from 0 to 80

tabulardata.d4science.org should appear in the right ganglia cluster soon.

I've removed ganglia from the other nodes. Is it possible that they were manually cloned from a smartgears worker node? It seems that they never appeared in the inventory host list.

#4 - Aug 04, 2015 04:26 PM - Andrea Dell'Amico

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

- % *Done* changed from 80 to 100