

OpenNebula - Feature #846

Scheduler approach on impossible results

09/29/2011 09:52 AM - Tino Vázquez

Status:	Closed	Start date:	09/29/2011
Priority:	Normal	Due date:	
Assignee:	Carlos Martín	% Done:	0%
Category:	Scheduler	Estimated time:	0.00 hour
Target version:	Release 4.0	Pull request:	
Resolution:	fixed		
Description			
<p>As the OpenNebula default matchmaking scheduler is an immediate lease scheduler, it should always return the allocated resources or an error due to it's impossibility. But, when the scheduler receives an expression which turns to an impossible result, the VM remains in pending state indefinitely, which is an approach related with best-effort schedulers.</p> <p>Specifically in our case, we want to deploy one Virtual Machine on each host of our cloud (for security reasons) and we don't want to obtain all information about the entire HostPool in each VM creation (scalability reasons).</p> <p>The problem arises when we achieve the entire set of hosts inside of our requirements (e.g using "HOSTNAME != \"some-host\" & \"HOSTNAME != \"another-host\"", etc), which means that don't exist other host to allocate the Virtual Machine.</p> <p>This approach also can be considered one of the reasons that contributed for other similar issues pointed on mailing list, as for example, on [one-users] VM creation stuck in "pending" state thread, where the same problems could be solved if the OpenNebula would changed the VM state to fail and logged the appropriate output.</p> <p>We believe that there are at least two main possibilities to circumvent this problem:</p> <ol style="list-style-type: none">1) Change the VM state to "fail" when an impossible result is achieved.2) Add a new state for the virtual machine to indicate when it achieve an impossible result (for example, an impossible state).			

Associated revisions

Revision 85e921de - 01/30/2013 11:56 AM - Carlos Martín

Feature #846: Minor changes in log messages

Revision cdb2a8a4 - 01/30/2013 04:49 PM - Carlos Martín

Feature #846: Add methods to update the VM template from the sched and log messages

Revision 10107cd6 - 01/30/2013 05:41 PM - Carlos Martín

Feature #846: Skip the scheduling if there aren't any pending VMs, and omit empty logs

Revision d51444dc - 01/31/2013 02:40 PM - Carlos Martín

Feature #846: Minor fix for req. syntax errors

Revision 1995b7a1 - 02/01/2013 01:23 AM - Ruben S. Montero

feature #846: Add more message types. Template is loaded in the VirtualMachineXML object. OpenNebula API calls are always made through VirtualMachinePoolXML.

History

#1 - 09/29/2011 09:53 AM - Tino Vázquez

This feature was suggested by VinÃ-cius Vielmo Cogo [vielmo@inf.ufsm.b](mailto:vielmo@inf.ufsm.br)

#2 - 10/04/2011 08:28 AM - Ruben S. Montero

- Target version deleted (Release 3.4)

#3 - 01/15/2013 09:33 AM - Ruben S. Montero

- Target version set to Release 4.0

#4 - 01/29/2013 04:48 PM - Carlos MartÃn

- Assignee set to Carlos MartÃn

#5 - 02/28/2013 03:01 AM - Ruben S. Montero

- Status changed from New to Closed

- Resolution set to fixed