

D4Science Infrastructure - Task #9956

Please slow down HAProxy/Nagios checks on Resource Registry

Oct 16, 2017 02:09 PM - Luca Frosini

Status:	Closed	Start date:	Oct 16, 2017
Priority:	Normal	Due date:	
Assignee:	_InfraScience Systems Engineer	% Done:	100%
Category:	System Application	Estimated time:	0.00 hour
Target version:	Nagios monitoring		
Infrastructure:	Production		
Description			
Looking at logs and accounting I found that the service receives 1,33 checks per second (one check every 0,75 second). The normal rate should be 1 check every 3 seconds. Can you please investigate it?			

History

#1 - Oct 16, 2017 02:49 PM - Andrea Dell'Amico

- Status changed from New to In Progress

The count is correct. haproxy checks every 3 seconds. But there are two haproxy instances, and each instance checks both http and https.

As each protocol receives, from the active haproxy, a check every 3 seconds and 4 failures are needed to declare a service dead. So, 12 seconds are needed at worst before disconnecting a non responsive server.

To me, 12 seconds is way too much. But if you can accept longer disservices, we can increase the check interval.

#2 - Oct 16, 2017 04:11 PM - Luca Frosini

The load balancer is contacted on by using https. The instances have to respond on http too because of the whn-manager. Can you please remove the http from load balancer?

#3 - Oct 16, 2017 05:02 PM - Andrea Dell'Amico

- Status changed from In Progress to Feedback

- % Done changed from 0 to 100

Done. The load balancer now balances the resource registry https accesses only.

#4 - Oct 16, 2017 06:39 PM - Luca Frosini

The requests are now about 40 per minutes on each instance (they were 80 per minute). This means that arrive one request every 1.5 second and not a request every 3 seconds.

If I correctly understand now 6 seconds are needed at worst before disconnecting a non responsive server.

Can we slow down to 1 request every 2 seconds which can be good compromise (8 seconds are needed at worst before disconnecting a non responsive server)

#5 - Oct 16, 2017 06:46 PM - Andrea Dell'Amico

Luca Frosini wrote:

The requests are now about 40 per minutes on each instance (they were 80 per minute). This means that arrive one request every 1.5 second and not a request every 3 seconds.

If I correctly understand now 6 seconds are needed at worst before disconnecting a non responsive server.

No, nothing changed: the check interval is always 3 seconds, and the worst case scenario is 12 seconds. The two haproxy are independent one from another, and each backend inside the same haproxy instance is independent too.

Can we slow down to 1 request every 2 seconds which can be good compromise (8 seconds are needed at worst before disconnecting a non responsive server)

That would move the worst case scenario to 16 seconds, not 8.

#6 - Oct 16, 2017 08:13 PM - Luca Frosini

- *Status changed from Feedback to Closed*

Ok I understood. I'm going to close the ticket