
Subject: HA for VE haproxy

Posted by [ktwalrus](#) on Sun, 25 Mar 2012 02:54:11 GMT

[View Forum Message](#) <> [Reply to Message](#)

I'm looking into adding High Availability for haproxy Virtual IP.

Let's say I have 4 HN's each with a VE running Centos 6 and haproxy for load balancing.

My research seems to indicate I might need to install and use Keepalived to make the haproxy VEs HA.

How do I best do this? My VEs all currently use venet0 for network access. Do I install Keepalived on each HN? Should the VIP be added to each VE so when Keepalived changes eth0 to route this IP to the VE on the HN?

I'm a bit confused as to how all this works. There is a Wiki page on the subject, but it seems half written doesn't go into how this all could work (just says to add NET_ADMIN capability to VE and add some kernel modules to the HN).

Subject: Re: HA for VE haproxy

Posted by [ktwalrus](#) on Fri, 14 Dec 2012 00:33:57 GMT

[View Forum Message](#) <> [Reply to Message](#)

I'm still interested in this even though I just put the haproxy/keepalived services on non-OpenVZ servers.

Is anyone running haproxy in two OpenVZ containers (that may be running on different HN's) and using keepalived/heartbeat for HA? Seems like the wiki should be updated on this subject. Must keepalived be installed in the HN's or can it be installed in the haproxy containers?
