
Subject: openvz vrrp (and ucarp)
Posted by [Meikes](#) on Fri, 26 Aug 2011 14:01:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hello,

I have running two VE, connected to the same bridge on the host.
Network communication is running well.

The problem occurs with program ucarp, version 1.5.1 under Debian 6. The IP protocol VRRP (0x70) does not work between VEs or VE and Host. If I sniff (tcpdump) inside the VE, I see the packages, but not on the host or the other VE.

If I change the IP protocol back to TCP (6) or UDP (0x11), the packages are seen on the host and also the other vz.

I used packETH as packet generator. Because I can reproduce the problem with the packet generator, it seems, it is not a problem from ucarp.

A "few" time ago, the same problem was reported but no one did answer.

forum.openvz.org/index.php?t=msg&goto=37736&&src_h=vrrp#msg_37736

Please help!

Kindly regards Meike!

Subject: Re: openvz vrrp (and ucarp)
Posted by [mose](#) on Sun, 26 Feb 2012 04:47:20 GMT
[View Forum Message](#) <> [Reply to Message](#)

After extensive testing, I have come to the same conclusion. veth is not passing uCarp's VRRP packets to CT0. The container is bridged to CT0's eth0. With tcpdump, I can see uCarp packets getting to eth0 in the container, but they don't come out the veth on the other side.

This is rather odd, because VRRP packets from the network make it across the bridge from CT0 to the container. I would think it would work both ways. I have confirmed that multicast is possible from the container with ssmping (which uses UDP), but nothing happens when the IP protocol is VRRP.

Subject: Re: openvz vrrp (and ucarp)

Posted by [mose](#) on Sun, 26 Feb 2012 22:23:11 GMT

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

I found some additional information that confirms our results. The only protocols allowed are essentially UDP, TCP and ICMP when using a stream or datagram socket. If a raw socket is opened and the protocol is IP, then other IP types could be used, such as VRRP.

So, this has nothing to do with multicasting, rather that the IP protocol type used by uCarp is VRRP. Since uCarp opens a datagram socket and not a raw socket, it will not work in this environment. Note that VRRP packets can be received by a container from the network as usual.

References:

forum.openvz.org/index.php?t=msg&goto=12877&bugzilla.openvz.org/show_bug.cgi?id=611

NB: It is quite annoying that I cannot post links.

Subject: Re: openvz vrrp (and ucarp)

Posted by [Meikes](#) on Wed, 29 Feb 2012 12:02:09 GMT

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

Hello,

it is not clear why, but "mac_filter=no" solved the problem....

see: (You cannot use links until you have posted more than 10 messages...)
"xxxx://bugzilla.openvz.org/show_bug.cgi?id=1967"

Kindly regards Meike

Subject: Re: openvz vrrp (and ucarp)

Posted by [mose](#) on Wed, 29 Feb 2012 22:11:46 GMT

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

That was very helpful. Setting mac_filter to off fixed the problem.

The documentation is now more clear. When mac_filter is turned on (the default), the container can only send packets with the MAC address that is assigned to the container. The container is not allowed to send packets with any other MAC address.

uCarp sends packets from a multicast MAC address by default, but those packets are not accepted by veth. You can have uCarp send from a broadcast MAC address, as well, but those packets are not accepted for the same reason. Setting mac_filter off allows the container to use whatever MAC address it wants.

Since this is a moderated forum, requiring 10 posts before allowing links seems more punitive than useful.
