
Subject: Re: Using multicast in virtual servers
Posted by [dev](#) on Wed, 26 Sep 2007 07:35:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

uname -a
(from both VEs if you VE <-> VE multicasting)

Thanks,
Kirill
P.S. plz check my another today email about veth code.

Peter Hinse wrote:

> Kirill Korotaev wrote:

>
>>Peter,
>>
>>Is the same setup working without openvz?
>>Have you used multicast before? Multicast is a bit complex to set up,
>>requires support from routers/switches etc., so this might well be not openvz-related.
>>But we setting up this test case right now to check ourselves.

>
>
> We use several JBoss clusters with multicast in our datacenters.

> Multicast itself works between the physical hosts:

>
> ssmping joined (S,G) = (192.168.198.54,232.43.211.234)
> pinging S from 192.168.198.53
> unicast from 192.168.198.54, seq=1 dist=0 time=1.513 ms
> multicast from 192.168.198.54, seq=1 dist=0 time=1.528 ms
> unicast from 192.168.198.54, seq=2 dist=0 time=0.107 ms
> multicast from 192.168.198.54, seq=2 dist=0 time=0.115 ms
> unicast from 192.168.198.54, seq=3 dist=0 time=0.097 ms
> multicast from 192.168.198.54, seq=3 dist=0 time=0.104 ms
> unicast from 192.168.198.54, seq=4 dist=0 time=0.113 ms
> multicast from 192.168.198.54, seq=4 dist=0 time=0.123 ms
> unicast from 192.168.198.54, seq=5 dist=0 time=0.102 ms
> multicast from 192.168.198.54, seq=5 dist=0 time=0.113 ms
> unicast from 192.168.198.54, seq=6 dist=0 time=0.119 ms
> multicast from 192.168.198.54, seq=6 dist=0 time=0.130 ms

>
> --- 192.168.198.54 statistics ---
> 6 packets transmitted, time 5276 ms
> unicast:
> 6 packets received, 0% packet loss
> rtt min/avg/max/std-dev = 0.097/0.341/1.513/0.524 ms
> multicast:
> 6 packets received, 0% packet loss since first mc packet (seq 1) recvd

> rtt min/avg/max/std-dev = 0.104/0.352/1.528/0.526 ms
>
> If I try the ssmpingd/ssmping between two virtual instances on two
> different hosts (or from one virtual instance to a physical host and vice
> versa):
>
> pinging S from 192.168.198.142
> unicast from 192.168.198.132, seq=1 dist=1 time=2627.063 ms
> unicast from 192.168.198.132, seq=2 dist=1 time=1626.828 ms
> unicast from 192.168.198.132, seq=3 dist=1 time=626.718 ms
> unicast from 192.168.198.132, seq=4 dist=1 time=0.100 ms
> unicast from 192.168.198.132, seq=5 dist=1 time=0.101 ms
> unicast from 192.168.198.132, seq=6 dist=1 time=0.150 ms
>
> --- 192.168.198.132 statistics ---
> 6 packets transmitted, time 5372 ms
> unicast:
> 6 packets received, 0% packet loss
> rtt min/avg/max/std-dev = 0.100/813.493/2627.063/997.511 ms
> multicast:
> 0 packets received, 100% packet loss
>
>
>
>> Can you please also provide a bit more information about your configuration
>> like whether you use bridge for veth-eth0 traffic bridging or routed networking,
>> any configuration options (including sysctl) you used/changed etc.?
>
>
> sysctl settings for the virtual interface on one of the host systems:
>
> net.ipv4.conf.veth1981320.promote_secondaries = 0
> net.ipv4.conf.veth1981320.force_igmp_version = 0
> net.ipv4.conf.veth1981320.disable_policy = 0
> net.ipv4.conf.veth1981320.disable_xfrm = 0
> net.ipv4.conf.veth1981320.arp_accept = 0
> net.ipv4.conf.veth1981320.arp_ignore = 0
> net.ipv4.conf.veth1981320.arp_announce = 0
> net.ipv4.conf.veth1981320.arp_filter = 0
> net.ipv4.conf.veth1981320.tag = 0
> net.ipv4.conf.veth1981320.log_martians = 0
> net.ipv4.conf.veth1981320.bootp_relay = 0
> net.ipv4.conf.veth1981320.medium_id = 0
> net.ipv4.conf.veth1981320.proxy_arp = 1
> net.ipv4.conf.veth1981320.accept_source_route = 1
> net.ipv4.conf.veth1981320.send_redirects = 1
> net.ipv4.conf.veth1981320.rp_filter = 0
> net.ipv4.conf.veth1981320.shared_media = 1

```
> net.ipv4.conf.veth1981320.secure_redirects = 1
> net.ipv4.conf.veth1981320.accept_redirects = 1
> net.ipv4.conf.veth1981320.mc_forwarding = 0
> net.ipv4.conf.veth1981320.forwarding = 1
>
>
> Network config on the VPS:
>
> eth0    Link encap:Ethernet HWaddr 00:0C:29:0A:3D:C6
>         inet addr:192.168.198.132 Bcast:192.168.198.255
>         Mask:255.255.255.0
>             inet6 addr: fe80::20c:29ff:fe0a:3dc6/64 Scope:Link
>                 UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
>                 RX packets:2855871 errors:0 dropped:0 overruns:0 frame:0
>                 TX packets:5495903 errors:0 dropped:0 overruns:0 carrier:0
>                 collisions:0 txqueuelen:0
>                 RX bytes:1453942778 (1.3 GiB) TX bytes:3721100335 (3.4 GiB)
>
> Kernel IP routing table
> Destination     Gateway         Genmask        Flags Metric Ref  Use
> Iface
> 192.168.198.0   0.0.0.0       255.255.255.0 U   0   0      0 eth0
> 169.254.0.0     0.0.0.0       255.255.0.0   U   0   0      0 eth0
> 224.0.0.0       0.0.0.0       240.0.0.0    U   0   0      0 eth0
>
> Any more information you need?
>
> Regards,
>
> Peter
>
```
