RHCS6: Debug and test multicast traffic between two hosts

Article Number: 205 | Rating: Unrated | Last Updated: Sun, Jun 3, 2018 9:52 AM

RHCS: Debug and test multicast traffic between two hosts

```
# Tested on RHEL 5 & 6


# Sometimes we may suspect of multicast traffic not working as
expected. In this case we
# may carry out following tests to figure out whether it is working
or not.


# On a RHEL 5, with <239.111.0.22> as multicast IP


# 'netstat -g' shows the interfaces' multicast group memberships

netstat -g
   IPv6/IPv4 Group Memberships
   Interface       RefCnt Group
   --------------- ------ ---------------------
   lo              1      all-systems.mcast.net
   eth2            2      all-systems.mcast.net
   eth3            1      239.111.0.22
   eth3            2      all-systems.mcast.net
   bond0           2      all-systems.mcast.net
```

```
   lo              1      ff02::1
   eth2            1      ff02::1:ff5b:352
   eth2            1      ff02::1
   eth3            1      ff02::1:ff5b:353
   eth3            1      ff02::1
   bond0           1      ff02::3:1
   bond0           1      ff02::1:ffc9:f168
   bond0           1      ff02::1



# 'netstat -s' shows a multicast packet counter that should increase
when traffic is
# received/sent

netstat -s | grep Mcast
    InMcastPkts: 378347
    OutMcastPkts: 230473

netstat -s | grep Mcast
    InMcastPkts: 378365
    OutMcastPkts: 230488



# 'tcpdump' shows the network traffic (eth3 being my cluster
interface)

tcpdump -i eth3 | grep 239.111.0.22
   tcpdump: verbose output suppressed, use -v or -vv for full
protocol decode
   listening on eth3, link-type EN10MB (Ethernet), capture size 96
bytes
   13:59:16.712278 IP myhost-priv.5149 > 239.111.0.22.netsupport:
UDP, length 118
   13:59:17.116244 IP myhost-priv.5149 > 239.111.0.22.netsupport:
UDP, length 118
   13:59:17.512239 IP myhost-priv.5149 > 239.111.0.22.netsupport:
```

```
UDP, length 118
   13:59:17.908238 IP myhost-priv.5149 > 239.111.0.22.netsupport:
UDP, length 118
   13:59:18.304221 IP myhost-priv.5149 > 239.111.0.22.netsupport:
UDP, length 118
   13:59:18.700217 IP myhost-priv.5149 > 239.111.0.22.netsupport:
UDP, length 118
   13:59:19.096197 IP myhost-priv.5149 > 239.111.0.22.netsupport:
UDP, length 118
   13:59:19.492195 IP myhost-priv.5149 > 239.111.0.22.netsupport:
UDP, length 118
   360 packets captured
   360 packets received by filter
   0 packets dropped by kernel




# On a RHEL 6, apart from the commands shown above, we can use
'omping'. It has to be
# started on all nodes indicating the IPs of the remote node(s) and
server's own IP as
# parameters:


myhostA:#> omping myhostA myhostB
   myhostA : waiting for response msg
   myhostA : waiting for response msg
   myhostA : waiting for response msg
   myhostA : waiting for response msg
   myhostA : waiting for response msg
   myhostA : joined (S,G) = (*, 232.43.211.234), pinging
   myhostA :   unicast, seq=1, size=69 bytes, dist=0, time=0.264ms
   myhostA : multicast, seq=1, size=69 bytes, dist=0, time=0.271ms
   myhostA :   unicast, seq=2, size=69 bytes, dist=0, time=0.312ms
   myhostA : multicast, seq=2, size=69 bytes, dist=0, time=0.320ms
   myhostA :   unicast, seq=3, size=69 bytes, dist=0, time=0.279ms
```

```
   myhostA : multicast, seq=3, size=69 bytes, dist=0, time=0.287ms

   myhostA :   unicast, xmt/rcv/%loss = 3/3/0%, min/avg/max/std-dev =
0.264/0.285/0.312/0.025
   myhostA : multicast, xmt/rcv/%loss = 3/3/0%, min/avg/max/std-dev =
0.271/0.293/0.320/0.025


myhostB:#> omping myhostA myhostB
   myhostB : waiting for response msg
   myhostB : joined (S,G) = (*, 232.43.211.234), pinging
   myhostB :   unicast, seq=1, size=69 bytes, dist=0, time=0.300ms
   myhostB : multicast, seq=1, size=69 bytes, dist=0, time=0.306ms
   myhostB :   unicast, seq=2, size=69 bytes, dist=0, time=0.325ms
   myhostB : multicast, seq=2, size=69 bytes, dist=0, time=0.331ms
   myhostB :   unicast, seq=3, size=69 bytes, dist=0, time=0.325ms
   myhostB : multicast, seq=3, size=69 bytes, dist=0, time=0.332ms
   myhostB :   unicast, seq=4, size=69 bytes, dist=0, time=0.353ms
   myhostB : multicast, seq=4, size=69 bytes, dist=0, time=0.359ms

   myhostB :   unicast, xmt/rcv/%loss = 4/4/0%, min/avg/max/std-dev =
0.300/0.326/0.353/0.022
   myhostB : multicast, xmt/rcv/%loss = 4/4/0%, min/avg/max/std-dev =
0.306/0.332/0.359/0.022




# Another idea could be making nodes answer multicast pings. In a
normal configuration
# when the multicast address is pinged by any node in the cluster
there is no response.
# By enabling multicast acknowledgements we will be able to receive a
response to our
# pings. If multicast is working well, all of the nodes should answer
the ping.
```

```
# To enable this functionality temporarily, run following command on
all nodes

sysctl -w net.ipv4.icmp_echo_ignore_broadcasts=0

# and test (For my cluster, formed by the nodes 192.168.100.101 and
192.168.100.102 with
# 239.111.0.22 as multicast address):

ping 239.111.0.22
   PING 239.111.0.22 (239.111.0.22) 56(84) bytes of data.
   64 bytes from 192.168.100.102: icmp_seq=1 ttl=64 time=0.027 ms
   64 bytes from 192.168.100.101: icmp_seq=1 ttl=64 time=0.334 ms
(DUP!)
   64 bytes from 192.168.100.102: icmp_seq=2 ttl=64 time=0.026 ms
   64 bytes from 192.168.100.101: icmp_seq=2 ttl=64 time=0.480 ms
(DUP!)
   64 bytes from 192.168.100.102: icmp_seq=3 ttl=64 time=0.029 ms
   64 bytes from 192.168.100.101: icmp_seq=3 ttl=64 time=0.309 ms
(DUP!)

   --- 239.111.0.22 ping statistics ---
   3 packets transmitted, 3 received, +3 duplicates, 0% packet loss,
time 2630ms
   rtt min/avg/max/mdev = 0.026/0.200/0.480/0.182 ms



# To make this change permanent add following line to
/etc/sysctl.conf

net.ipv4.icmp_echo_ignore_broadcasts = 0

# and load the new setting

systcl -p
```