RHEL: Route network packets to go out via the same interface they came in

Article Number: 131 | Rating: Unrated | Last Updated: Sat, Jun 2, 2018 8:52 AM

RHEL: Route network packets to go out via the same interface they came in

```
# Tested on RHEL 6



# When working on a server with several network interfaces, if we don't define any specific

# VLAN routing, all the outgoing traffic will usually go through the default interface.



# On servers connected to many different VLANs, and with special requirements regarding

# the traffic balancing on the physical interfaces, such may be the case of a backup server,

# this could be a laborious issue to manage.
```

```
# I have one bacukp server with two network interfaces: the first
one, bond0, connected to

# the administrative network and the other one, bond1, linked to the
backup network and

# running a service on a virtual IP that will be used by clients to
send their data:



[root@mybckserver ~]# ifconfig | egrep "Link|inet add"

bond0  Link encap:Ethernet  HWaddr 24:6E:G6:H7:99:14

  inet addr:10.69.961.124  Bcast:10.69.961.255  Mask:255.255.255.0

bond1  Link encap:Ethernet  HWaddr A0:36:MF:C8:DC:88

  inet addr:10.256.11.117  Bcast:10.256.11.255  Mask:255.255.255.0

bond1:0  Link encap:Ethernet  HWaddr A0:36:MF:C8:DC:88

  inet addr:10.256.11.118  Bcast:10.256.11.255  Mask:255.255.255.0

eth0  Link encap:Ethernet  HWaddr 24:6E:G6:H7:99:14

eth1  Link encap:Ethernet  HWaddr 24:6E:G6:H7:99:15

eth2  Link encap:Ethernet  HWaddr 24:6E:G6:H7:99:14

eth3  Link encap:Ethernet  HWaddr 24:6E:G6:H7:99:17

eth4  Link encap:Ethernet  HWaddr A0:36:MF:C8:DC:88

eth6  Link encap:Ethernet  HWaddr A0:36:MF:C8:DC:88

lo  Link encap:Local Loopback
```

```
  inet addr:127.0.0.1  Mask:255.0.0.0




# To avoid overloading the public administrative network, I would
like to configure the

# backup interface so all the requests are answered over this
interface, this is, I'll force

# all outgoing packets to go out via the interface they came in.



# With my current configuration, all the ping requests made to my
backup service, "bck-srv",

# are answered via bond0 instead of bond1:



[root@client01 ~]# ping -c 2 bck-srv

PING bck-srv.syscookbook.mydomain.org (10.256.11.118) 56(84) bytes of
data.

64 bytes from bck-srv.syscookbook.mydomain.org (10.256.11.118):
icmp_seq=1 ttl=63 time=0.301 ms

64 bytes from bck-srv.syscookbook.mydomain.org (10.256.11.118):
icmp_seq=2 ttl=63 time=0.333 ms



--- bck-srv.syscookbook.mydomain.org ping statistics ---
```

```
2 packets transmitted, 2 received, 0% packet loss, time 999ms

rtt min/avg/max/mdev = 0.301/0.317/0.333/0.016 ms



# Incoming traffic arriving on bond1...



root@mybckserver:/root#> tcpdump -i bond1 host client01

tcpdump: verbose output suppressed, use -v or -vv for full protocol
decode

listening on bond1, link-type EN10MB (Ethernet), capture size 65535
bytes

13:06:26.359950 IP client01.syscookbook.mydomain.org > bck-
srv.syscookbook.mydomain.org: ICMP echo request, id 14191, seq 1,
length 64

13:06:27.359453 IP client01.syscookbook.mydomain.org > bck-
srv.syscookbook.mydomain.org: ICMP echo request, id 14191, seq 2,
length 64



2 packets captured

3 packets received by filter

0 packets dropped by kernel



# ...was answered using bond0:
```

```
root@mybckserver:/root#> tcpdump -i bond0 host client01

tcpdump: verbose output suppressed, use -v or -vv for full protocol
decode

listening on bond0, link-type EN10MB (Ethernet), capture size 65535
bytes

13:06:26.360003 IP bck-srv.syscookbook.mydomain.org >
client01.syscookbook.mydomain.org: ICMP echo reply, id 14191, seq 1,
length 64

13:06:27.359477 IP bck-srv.syscookbook.mydomain.org >
client01.syscookbook.mydomain.org: ICMP echo reply, id 14191, seq 2,
length 64



2 packets captured

13 packets received by filter

0 packets dropped by kernel




# This was the default routing table:



root@mybckserver:/root#> route -n

Kernel IP routing table
```

```
Destination  Gateway  Genmask  Flags Metric Ref  Use Iface

10.256.11.0  0.0.0.0  255.255.255.0  U  0  0  0 bond1

10.69.961.0  0.0.0.0  255.255.255.0  U  0  0  0 bond0

169.254.0.0  0.0.0.0  255.255.0.0  U  1010  0  0 bond0

169.254.0.0  0.0.0.0  255.255.0.0  U  1011  0  0 bond1

0.0.0.0  10.69.961.1  0.0.0.0  UG  0  0  0 bond0



# No special route or rule declared:



root@mybckserver:/root#> ls -lrt /etc/sysconfig/network-scripts/rule*

ls: cannot access /etc/sysconfig/network-scripts/rule*: No such file
or directory



root@mybckserver:/root#> ls -lrt /etc/sysconfig/network-
scripts/route*

ls: cannot access /etc/sysconfig/network-scripts/route*: No such file
or directory



# At this point, to redirect all the backup traffic through the
```

```
dedicated interface, I

# would need to define new routes for each VLAN connecting to my
server for backups

# purposes.



# Should I do this, on one hand, in the long term I would have an
enormous routing table,

# not easy to manage, and on the other hand I would run the risk of
forgetting adding new

# VLANs to the routing table so I thought that it would be better to
redirect ALL the

# replies to the requests received on bond1 through this interface.




# First, I create a new routing table dedicated to bond1:



root@mybckserver:/root#> ip route add 10.256.11.0/24 dev bond1 table
1

root@mybckserver:/root#> ip route add default via 10.256.11.1 dev
bond1 table 1



# And then I add some rules to link bond1 to the previously created
table and process all
```

```
# the traffic received on the corresponding IPs by this new routing
table:



root@mybckserver:/root#> ip rule add iif bond1 table 1



root@mybckserver:/root#> ip rule add from 10.256.11.117 table 1

root@mybckserver:/root#> ip rule add from 10.256.11.118 table 1




# This is the my configuration:



root@mybckserver:/root#> ls -lrt /etc/sysconfig/network-scripts/rule*

-rw-r--r-- 1 root root 70 Feb 20 13:17 /etc/sysconfig/network-
scripts/rule-bond1



root@mybckserver:/root#> ls -lrt /etc/sysconfig/network-
scripts/route*

-rw-r--r-- 1 root root 72 Feb 20 13:17 /etc/sysconfig/network-
scripts/route-bond1



root@mybckserver:/root#> cat /etc/sysconfig/network-scripts/route-
bond1
```

```
10.256.11.0/24 dev bond1 table 1

default via 10.256.11.1 dev bond1 table 1



root@mybckserver:/root#> cat /etc/sysconfig/network-scripts/rule-bond1

iif bond1 table 1

from 10.256.11.117 table 1

from 10.256.11.118 table 1




# I'll restart the network service to verify that everything is ok with this configuration



root@mybckserver:/root#> service network restart




# And, as a result, all the traffic received on bond1 is using the default gateway of the

# new routing table and thus all the backup traffic will go through the dedicated interface,

# as expected ("dsmc" makes part of the commercial suite I'm using for backups):
```

```
[root@client01 ~]# dsmc

IBM Tivoli Storage Manager

Command Line Backup-Archive Client Interface

  Client Version 7, Release 1, Level 6.4

  Client date/time: 02/20/2018 13:14:06

(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights
Reserved.



Node Name: client01

Session established with server MYSERVER: Linux/x86_64

  Server Version 7, Release 1, Level 8.0

  Server date/time: 02/20/2018 13:14:06  Last access: 02/14/2018
15:56:30



tsm> quit



# Incoming traffic arriving on bond1...



root@mybckserver:/root#> tcpdump -i bond1 host client01
```

```
tcpdump: verbose output suppressed, use -v or -vv for full protocol
decode

listening on bond1, link-type EN10MB (Ethernet), capture size 65535
bytes

[...]

13:14:06.681180 IP client01.syscookbook.mydomain.org.40916 > bck-
srv.syscookbook.mydomain.org.imtc-mcs: Flags [.], ack 7978, win 169,
length 0

13:14:10.899841 IP client01.syscookbook.mydomain.org.40916 > bck-
srv.syscookbook.mydomain.org.imtc-mcs: Flags [P.], seq 572:576, ack
7978, win 169, length 4

13:14:10.900082 IP client01.syscookbook.mydomain.org.40916 > bck-
srv.syscookbook.mydomain.org.imtc-mcs: Flags [F.], seq 576, ack 7978,
win 169, length 0

13:14:10.913489 IP bck-srv.syscookbook.mydomain.org.imtc-mcs >
client01.syscookbook.mydomain.org.40916: Flags [F.], seq 7978, ack
577, win 149, length 0

13:14:10.913812 IP client01.syscookbook.mydomain.org.40916 > bck-
srv.syscookbook.mydomain.org.imtc-mcs: Flags [.], ack 7979, win 169,
length 0



28 packets captured

28 packets received by filter

0 packets dropped by kernel
```

```
# ...not being replied anymore via bond0:


root@mybckserver:/root#> tcpdump -i bond0 host client01

tcpdump: verbose output suppressed, use -v or -vv for full protocol
decode

listening on bond0, link-type EN10MB (Ethernet), capture size 65535
bytes


0 packets captured

10 packets received by filter

0 packets dropped by kernel




# ...now I can't see the new configuration for bond1 on the default routing table


root@mybckserver:/root#> route -n

Kernel IP routing table

Destination  Gateway  Genmask  Flags Metric Ref  Use Iface

10.256.11.0  0.0.0.0  255.255.255.0  U  0  0  0 bond1
```

```
10.69.961.0  0.0.0.0  255.255.255.0  U  0  0  0 bond0

169.254.0.0  0.0.0.0  255.255.0.0  U  1010  0  0 bond0

169.254.0.0  0.0.0.0  255.255.0.0  U  1011  0  0 bond1

0.0.0.0  10.69.961.1  0.0.0.0  UG  0  0  0 bond0




# For that I have to look directly in the new routing table:


root@mybckserver:/root#> ip route show table 1

10.256.11.0/24 dev bond1  scope link

default via 10.256.11.1 dev bond1
```