

In this tutorial I took Jake Edge's Namespaces in operation (<https://lwn.net/Articles/580893/>), and added some stuff.

1. Two VMs: lin70 and lin71
2. Addresses lin70: 192.168.4.240
lin71: 192.168.4.241
3. Login to lin71
4. Check connectivity between the two VMs

```
lin71 #ping -w 1 192.168.4.240
PING 192.168.4.240 (192.168.4.240) 56(84) bytes of data.
64 bytes from 192.168.4.240: icmp_seq=1 ttl=64 time=0.137 ms
64 bytes from 192.168.4.240: icmp_seq=2 ttl=64 time=0.120 ms
```

5. Setup network namespace netns1

```
lin71 #ip netns add netns1
lin71 #ip netns exec netns1 ip link list
1: lo: <LOOPBACK> mtu 65536 qdisc noop state DOWN mode DEFAULT
qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00

lin71 #ip netns exec netns1 ping -w1 127.0.0.1
PING 127.0.0.1 (127.0.0.1) 56(84) bytes of data.
64 bytes from 127.0.0.1: icmp_seq=1 ttl=64 time=0.130 ms
64 bytes from 127.0.0.1: icmp_seq=2 ttl=64 time=0.063 ms
```

6.

```
lin71 #ip link add veth0 type veth peer name veth1
lin71 #ip link set veth1 netns netns1
lin71 #ip netns exec netns1 ifconfig veth1 10.1.1.1/24 up
lin71 #ifconfig veth0 10.1.1.2/24 up

lin71 #ping -w 1 10.1.1.1
PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=0.074 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.068 ms

lin71 #ip netns exec netns1 ping -w 1 10.1.1.2
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data.
64 bytes from 10.1.1.2: icmp_seq=1 ttl=64 time=0.104 ms
64 bytes from 10.1.1.2: icmp_seq=2 ttl=64 time=0.048 ms
```

Setup routing in netns1

```
lin71 #ip netns exec netns1 route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
10.1.1.0 0.0.0.0 255.255.255.0 U 0 0 0 veth1

lin71 #ip netns exec netns1 ping -w 1 192.168.4.241
connect: Network is unreachable
```

```
lin71 #ip netns exec netns1 ip route add default via 10.1.1.2

lin71 #ip netns exec netns1 ping -w 1 192.168.4.241
PING 192.168.4.241 (192.168.4.241) 56(84) bytes of data.
64 bytes from 192.168.4.241: icmp_seq=1 ttl=64 time=0.070 ms
64 bytes from 192.168.4.241: icmp_seq=2 ttl=64 time=0.053 ms
```

7. Check connectivity with lin70

```
lin71 #ip netns exec netns1 ping -w 1 192.168.4.240
PING 192.168.4.240 (192.168.4.240) 56(84) bytes of data.

--- 192.168.4.240 ping statistics ---
2 packets transmitted, 0 received, 100% packet loss, time 999ms
```

lin70 does not have a route back to 10.1.1.0
Add a routing entry on lin70

```
lin70 #ip route add 10.1.1.0/24 via 192.168.4.241
lin70 #netstat -rn
Kernel IP routing table
Destination Gateway Genmask Flags MSS Window irtt
Iface
0.0.0.0 192.168.4.1 0.0.0.0 UG 0 0 0 ens192
10.1.1.0 192.168.4.241 255.255.255.0 UG 0 0 0 ens192
192.168.4.0 0.0.0.0 255.255.255.0 U 0 0 0 ens192
```

8. Check connectivity again

```
9. lin71 #ip netns exec netns1 ping -w 1 192.168.4.240
PING 192.168.4.240 (192.168.4.240) 56(84) bytes of data.
64 bytes from 192.168.4.240: icmp_seq=1 ttl=63 time=0.350 ms
```

Run a bash shell in netns1

```
lin71 #ip netns exec netns1 unshare --fork --pid --mount-proc bash
lin71 #ps aux
```

| USER | PID | %CPU | %MEM | VSZ | RSS | TTY | STAT | START | TIME | COMMAND |
|------|-----|------|------|--------|------|-------|------|-------|------|---------|
| root | 1 | 0.0 | 0.1 | 115512 | 1988 | pts/0 | S | 08:20 | 0:00 | bash |
| root | 2 | 0.0 | 0.0 | 139500 | 1628 | pts/0 | R+ | 08:21 | 0:00 | ps aux |

10. Check connectivity again

```
lin71 #ping -w 1 192.168.4.240
PING 192.168.4.240 (192.168.4.240) 56(84) bytes of data.
64 bytes from 192.168.4.240: icmp_seq=1 ttl=63 time=0.160 ms
64 bytes from 192.168.4.240: icmp_seq=2 ttl=63 time=0.254 ms
```