Configuring Multihome Networking on VM Hosts

Multihome networking on SUSE Linux Enterprise Server for NVIDIA configures multiple network interfaces on a virtual machine host using the strong host model to enhance routing isolation and security. By restricting Address Resolution Protocol (ARP) responses and implementing reverse path filtering, multihoming prevents cross-interface IP spoofing, eliminates ARP flux and ensures reliable network traffic segregation across complex virtualized environments.

Prerequisites

- Access to a machine that serves as the VM host
- Basic understanding of networking and IP addresses

Publication Date: 05 Nov 2026

Contents

1 Network topology 3

2 Checking the VM host network configuration 3

3 Checking the VM network configuration 5

4 IP **`sysctl`** settings for the multihome host 8

5 Default behavior 11

6 Multihome setup 19

7 Legal Notice 25

A GNU Free Documentation License 26

1 Network topology

The network topology in this setup includes three virtual machines: `vm7`, `vm8`, `vm9` and one uplink router `gw1`.

- The `vm7` VM acts as a router between an uplink router `gw1` and the external network (192.168.7.0/24) of `vm8` and `vm9`. The configuration of the `vm7 uplink` is not relevant for this setup.
- The `vm8` and `vm9` have an uplink in the external network (192.168.7.0/24) with `vm7` as a router and are connected to two private networks (192.168.8.0/24 and 192.168.9.0/24).

```
   gw1 (1)
             |
          ---+----       net1: 192.168.1.0/24 on eth1/enp1s0
             |           [uplink network: not relevant here]
           vm7 (107)
             |
     --+-----+-----+--   net7: 192.168.7.0/24 on eth7/enp7s0
       |           |
(108) vm8         vm9 (109)
      | |         | |
      | +---------+ |    net8: 192.168.8.0/24 on eth8/enp8s0
      +-------------+    net9: 192.168.9.0/24 on eth9/enp9s0
```

2 Checking the VM host network configuration

The following procedure outlines the steps for setting up a VM host.

The VM host has four bridges involved in the setup.

When the VM host is on a productive network with IP connectivity independent of the test networks, there is no IP configured on any of these bridges, including `net1`. However, you can use the same network service as the host, such as DHCP on `net1`.

The `net1` bridge contains an uplink port `eth1` providing access to an external test/LAN network for the `vm7` router VM.

PROCEDURE 1: SETTING UP VM HOST

1. Check the configuration of the `ifcfg-net1` bridge:

```
$ cat /etc/sysconfig/network/ifcfg-net1

  STARTMODE='auto'
```

```
  BOOTPROTO='none'
  LINK_REQUIRED=no
  BRIDGE='yes'
  BRIDGE_STP='off'
  BRIDGE_FORWARDDELAY='0'
  BRIDGE_PORTS='eth1'
```

2. The `net7`, `net8` and `net9` are host-only bridges that have only dynamic VM ports, specific and relevant for this setup. To view the `net7` network settings, run:

```
$ cat /etc/sysconfig/network/ifcfg-net7

  STARTMODE='auto'
  BOOTPROTO='none'
  LINK_REQUIRED=no
  LLADDR=66:00:00:00:00:07
  BRIDGE='yes'
  BRIDGE_STP='off'
  BRIDGE_FORWARDDELAY='0'
  BRIDGE_PORTS=''
```

3. View the `ifcfg-net8` network settings:

```
$ cat /etc/sysconfig/network/ifcfg-net8

  STARTMODE='auto'
  BOOTPROTO='none'
  LINK_REQUIRED=no
  LLADDR=66:00:00:00:00:08
  BRIDGE='yes'
  BRIDGE_STP='off'
  BRIDGE_FORWARDDELAY='0'
  BRIDGE_PORTS=''
```

4. View the `ifcfg-net9` network settings:

```
$ cat /etc/sysconfig/network/ifcfg-net9

  STARTMODE='auto'
  BOOTPROTO='none'
  LINK_REQUIRED=no
  LLADDR=66:00:00:00:00:09
  BRIDGE='yes'
  BRIDGE_STP='off'
  BRIDGE_FORWARDDELAY='0'
  BRIDGE_PORTS=''
```

3 Checking the VM network configuration

The following procedure outlines how the network of a VM host must be configured.

Configuration of the VM network shown in this document uses a so-called tweaked assignment of the IP, MAC and interface names. Every packet uses identical VM identification, just like packet captures. The following list shows the IP addresses for VMs on each network:

- net1
  vm7: no relevant uplink
- net7
  - vm7: 192.168.7.107 52:54:00:00:07:07 (eth7/enp7s0)
  - vm8: 192.168.7.108 52:54:00:00:07:08 (eth7/enp7s0)
  - vm9: 192.168.7.109 52:54:00:00:07:09 (eth7/enp7s0)
- net8
  - vm8: 192.168.8.108 52:54:00:00:08:08 (eth8/enp8s0)
  - vm9: 192.168.8.109 52:54:00:00:08:09 (eth8/enp8s0)
- net9
  - vm8: 192.168.9.108 52:54:00:00:09:08 (eth9/enp9s0)
  - vm9: 192.168.9.109 52:54:00:00:09:09 (eth9/enp9s0)

PROCEDURE 2: CONFIGURE VIRTUAL MACHINES

1. Set up vm7.
   This setup uses `SLES 15 SP7`, with the VM acting as a router: it connects to the external network through the `eth1` uplink (for example, assigned to an external firewalld zone for masquerading) and to the `common` test network `net7`, which serves as the uplink for the other VMs. `vm7` has no connection to the `private`, `net8` and `net9` networks.

   a. View the IP address:

```
ip a s
[...]
2: eth1: >BROADCAST,MULTICAST,UP,LOWER_UP< mtu 1500 qdisc pfifo_fast state UP
 group default qlen 1000
```

```
    link/ether 52:54:00:00:01:07 brd ff:ff:ff:ff:ff:ff
    altname enp1s0
    inet 192.168.1.107/24 brd 192.168.1.255 scope global eth1
       valid_lft forever preferred_lft forever
3: eth7: >BROADCAST,MULTICAST,UP,LOWER_UP< mtu 1500 qdisc pfifo_fast state UP
 group default qlen 1000
    link/ether 52:54:00:00:07:07 brd ff:ff:ff:ff:ff:ff
    altname enp7s0
    inet 192.168.7.107/24 brd 192.168.7.255 scope global eth7
       valid_lft forever preferred_lft forever
```

b. View the routing table:

```
default via 192.168.1.1 dev eth1
192.168.1.0/24 dev eth1 proto kernel scope link src 192.168.1.107
192.168.7.0/24 dev eth7 proto kernel scope link src 192.168.7.107
```

c. In `/etc/sysctl.d/90-network.conf`, set the following value to `1` to ensure this machine acts as a router:

```
net.ipv4.conf.all.forwarding = 1
```

2. Set up `vm8`.

   This configuration runs on `SLES 15 SP7`, using `eth7` as the network uplink and including interfaces for two private networks: `net8` and `net9`.

   a. View the IP address:

```
>
  ip a s
3: eth7: >BROADCAST,MULTICAST,UP,LOWER_UP< mtu 1500 qdisc pfifo_fast state UP
 group default qlen 1000
    link/ether 52:54:00:00:07:08 brd ff:ff:ff:ff:ff:ff
    altname enp7s0
    inet 192.168.7.108/24 brd 192.168.7.255 scope global eth7
       valid_lft forever preferred_lft forever
4: eth8: >BROADCAST,MULTICAST,UP,LOWER_UP< mtu 1500 qdisc pfifo_fast state UP
 group default qlen 1000
    link/ether 52:54:00:00:08:08 brd ff:ff:ff:ff:ff:ff
    altname enp8s0
    inet 192.168.8.108/24 brd 192.168.8.255 scope global eth8
       valid_lft forever preferred_lft forever
5: eth9: >BROADCAST,MULTICAST,UP,LOWER_UP< mtu 1500 qdisc pfifo_fast state UP
 group default qlen 1000
    link/ether 52:54:00:00:09:08 brd ff:ff:ff:ff:ff:ff
    altname enp9s0
```

```
    inet 192.168.9.108/24 brd 192.168.9.255 scope global eth9
       valid_lft forever preferred_lft forever
```

b. View the network routes

```
> ip r s
      default via 192.168.7.107 dev eth7
192.168.7.0/24 dev eth7 proto kernel scope link src 192.168.7.108
192.168.8.0/24 dev eth8 proto kernel scope link src 192.168.8.108
192.168.9.0/24 dev eth9 proto kernel scope link src 192.168.9.108
```

c. Set the following value in `/etc/sysctl.d/90-network.conf` to `0` to ensure that the machine acts as a host.

```
net.ipv4.conf.all.forwarding = 0
```

3. Set up `vm9`.

   This configuration runs on `SLES 16.0`, using `eth7` as the network uplink and connecting to two private networks: `net8` and `net9`.

   a. View the IP address:

```
> ip a s
3: enp7s0: >BROADCAST,MULTICAST,UP,LOWER_UP< mtu 1500 qdisc pfifo_fast state
 UP group default qlen 1000
    link/ether 52:54:00:00:07:09 brd ff:ff:ff:ff:ff:ff
    altname enx525400a8076d
    inet 192.168.7.109/24 brd 192.168.7.255 scope global noprefixroute enp7s0
       valid_lft forever preferred_lft forever
4: enp8s0: >BROADCAST,MULTICAST,UP,LOWER_UP< mtu 1500 qdisc pfifo_fast state
 UP group default qlen 1000
    link/ether 52:54:00:00:08:09 brd ff:ff:ff:ff:ff:ff
    altname enx525400a8086d
    inet 192.168.8.109/24 brd 192.168.8.255 scope global noprefixroute enp8s0
       valid_lft forever preferred_lft forever
5: enp9s0: >BROADCAST,MULTICAST,UP,LOWER_UP< mtu 1500 qdisc pfifo_fast state
 UP group default qlen 1000
    link/ether 52:54:00:00:09:09 brd ff:ff:ff:ff:ff:ff
    altname enx525400a8096d
    inet 192.168.9.109/24 brd 192.168.9.255 scope global noprefixroute enp9s0
       valid_lft forever preferred_lft forever
```

   b. View the network routes:

```
>
```

```
ip r s

default via 192.168.7.107 dev enp7s0 proto static metric 100
192.168.7.0/24 dev enp7s0 proto kernel scope link src 192.168.7.109 metric
 100
192.168.8.0/24 dev enp8s0 proto kernel scope link src 192.168.8.109 metric
 101
192.168.9.0/24 dev enp9s0 proto kernel scope link src 192.168.9.109 metric 102
```

c. Set the following value in `/etc/sysctl.d/90-network.conf` to `0` to ensure that the machine acts as a host.

```
net.ipv4.conf.all.forwarding = 0
```

4 IP **sysctl** settings for the multihome host

System-level networking parameters in Linux are stored in `/proc/sys/net/` and can be managed using the **sysctl** command or by modifying configuration files.

4.1 The **sysctl** command variables

The **sysctl** variables offered by the Linux kernel are documented here (https://www.kernel.org/doc/Documentation/networking/ip-sysctl.txt).

The following `net.ipv4.conf.all.arp_*` variables affect the `arp` behavior on a multihome host.

`arp_announce = 0`

- Define different restriction levels for announcing the local source IP address from IP packets in ARP requests sent on the interface:

  1. `0` (default) Use any local address, configured on any interface.

  2. `1` Try to avoid local addresses that are not in the target's subnet for this interface. This mode is useful when target hosts reachable via this interface require the source IP address in ARP requests to be part of their logical network configured on the receiving interface. When gen-

erating the request, all subnets that include the target IP are checked. If such a subnet exists, the source address is preserved. Otherwise, the source address is selected according to the rules for level 2.

3. `2` Always use the best local address for this target. In this mode, the source address in the IP packet is ignored, and the preferred local address is selected from the outgoing interface that includes the target IP address. If no suitable address is found, the first local address on the outgoing interface or other interfaces is used, hoping to receive a reply.

The max value from `conf/{all,interface}/arp_announce` is used.
Increasing the restriction level gives a better chance of receiving an answer from the resolved target, while decreasing the level announces more valid sender information.

- `aarp_ignore = 0`

  Define different modes for sending replies in response to received ARP requests that resolve local target IP addresses:

  1. `0` (default) Reply to requests for any local target IP address, configured on any interface.
  2. `1` Reply only if the target IP address is configured on the incoming interface.
  3. `2` Reply only if the target IP address is configured on the incoming interface and both it and the sender's IP address are in the same subnet on this interface.
  4. `3`Do not reply to requests for local addresses configured with scope host; only resolutions for global and link addresses are replied.
  5. `4–7` Reserved.
  6. `8`Do not reply to requests for any local address.

  The max value from `conf/{all,interface}/arp_ignore` is used when ARP requests are received on the {interface}.

- `arp_filter = 0`

Controls ARP response behavior for interfaces on the same subnet:

1. `0` (default) The kernel can respond to ARP requests with addresses from other interfaces. This increases the chance of successful communication. IP addresses are owned by the complete host on Linux, not by particular interfaces. In more complex setups like load-balancing, this behavior can cause problems.

2. `1` Each interface responds only if the kernel would route the packet through it. Requires source-based routing and ensures ARP replies are interface-specific.

`arp_filter` for the interface will be enabled if at least one of `conf/{all,interface}/arp_filter` is set to TRUE; otherwise it is disabled.

Additionally, the `net.ipv4.conf.all.rp_filter` settings permit filtering of L3 / IPv4 packets using reverse routing table lookups:

- `rp_filter = 0`
  Modes for source validation:

  1. `0` No source validation.

  2. `1` Strict mode (RFC3704). Each incoming packet is tested against the FIB. If the interface is not the best reverse path, the packet check fails, and the packet is discarded.

  3. `2` Loose mode (RFC3704). Each incoming packet's source address is tested against the FIB, and if it is not reachable via any interface, the packet check fails.

  Current recommended practice in RFC3704 is to enable strict mode to prevent IP spoofing from DDoS attacks. If using asymmetric routing or other complicated routing, loose mode is recommended.
  The max value from `conf/{all,interface}/rp_filter` is used when doing source validation on the interface.
  The default value is `0`. Note that some distributions enable it in startup scripts.

On SUSE systems, `net.ipv4.conf.all.rp_filter = 2` (RFC3704 Loose Reverse Path) is enabled by default instead of the kernel default to disable it (`rp_filter = 0`). Unlike the strict `rp_filter = 1` mode, `rp_filter = 2` permits asymmetric routing, using the same subnet on multiple interfaces (considers `arp_announce=0`, `arp_ignore=0` and `arp_filter=0` defaults).

The `net.ipv4.conf.all.log_martians` enables logging of packets filtered/dropped by `rp_filter`:

```
log_martians = 0
```

This logs packets with impossible addresses to the kernel log. The `log_martians` setting for an interface will be enabled if at least one of `conf/{all,interface}/log_martians` is set to `TRUE`. Otherwise, it will be disabled.

5 Default behavior

Linux uses a weak host model (https://en.wikipedia.org/wiki/Host_model) as default, configurable via IP-Sysctl settings offered by the kernel.

In this model, an IP address belongs to the machine and is not restricted to being used only on the interface it is assigned to. The machine can use/associate any interface MAC address with any IP address, regardless of the interface.

The following procedures describe the default behavior using vm7, vm8, and vm9.

PROCEDURE 3: DEFAULT BEHAVIOR

1. vm7

   This VM acts as a router in `net7` and even though it is not connected to `net8` and `net9`, it is still involved in the ARP resolution process.

   The following ping commands illustrate the default behavior:

   ```
   vm7:~ ping -c 1 192.168.7.108: reply (regular)
   ```

   Regular pinging of an IP on a directly connected network (network route) triggers a routing table lookup that selects the `eth7` interface and the preferred source from the route (src 192.168.7.107):

   ```
   192.168.7.0/24 dev eth7 proto kernel scope link src 192.168.7.107
   ```

   Output:

   ```
     PING 192.168.7.108 (192.168.7.108) 56(84) bytes of data.
     64 bytes from 192.168.7.108: icmp_seq=1 ttl=64 time=0.446 ms

     --- 192.168.7.108 ping statistics ---
     1 packets transmitted, 1 received, 0% packet loss, time 0ms
     rtt min/avg/max/mdev = 0.446/0.446/0.446/0.000 ms

   - capture (net7):

   No.     Time           Source                Destination           Protocol Length
    Info
   ```

```
       1 0.000000000    52:54:00:00:07:07     Broadcast             ARP      42
Who has 192.168.7.108? Tell 192.168.7.107
       2 0.000178399    52:54:00:00:07:08     52:54:00:00:07:07     ARP      42
192.168.7.108 is at 52:54:00:00:07:08
       3 0.000261527    192.168.7.107         192.168.7.108         ICMP     98
Echo (ping) request  id=0x0001, seq=1/256, ttl=64 (reply in 4)
       4 0.000334225    192.168.7.108         192.168.7.107         ICMP     98
Echo (ping) reply    id=0x0001, seq=1/256, ttl=64 (request in 3)
       5 5.180898743    52:54:00:00:07:08     52:54:00:00:07:07     ARP      42
Who has 192.168.7.107? Tell 192.168.7.108
       6 5.181116066    52:54:00:00:07:07     52:54:00:00:07:08     ARP      42
192.168.7.107 is at 52:54:00:00:07:07
```

```
vm7:~ ping -c 1 192.168.7.109: reply (regular)
```

Same as `vm7:~ ping -c 1 192.168.7.108 (#vm7--ping--c-1-1921687108-reply)`.

```
vm7:~ ping -c 1 192.168.8.108: error (regular)
```

Regular pinging of an IP *outside* of a directly connected network causes the packet to be sent via the gateway (default) route. This uses the `eth1` interface with the source IP (`192.168.1.107`) from the preferred source in the route used to reach the gateway:

```
default via 192.168.1.1 dev eth1 192.168.1.0/24 dev eth1 proto kernel scope link src
 192.168.1.107
```

The result depends on the configuration of the uplink routers. If there is a `kind router` sending an `unreachable` ICMP message, for example, because of an `unreachable` route (on Linux, to avoid routing the packets to the Internet):

```
ip r s type unreachable
unreachable 192.168.0.0/16
```

or dropping the packets (because of a `blackhole` route entry or a firewall).
Output:

```
 PING 192.168.8.108 (192.168.8.108) 56(84) bytes of data.

  --- 192.168.8.108 ping statistics ---
  1 packets transmitted, 0 received, 100% packet loss, time 0ms

  or:

  PING 192.168.8.108 (192.168.8.108) 56(84) bytes of data.
  From 192.168.1.1 icmp_seq=1 Destination Host Unreachable
```

```
  --- 192.168.8.108 ping statistics ---
  1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

- capture (net1):
  - packet capture on net1 incl. unreachable message (if any):

  No.     Time           Source                Destination           Protocol Length
 Info

      131 1995.230774    192.168.1.107         192.168.8.108         ICMP     104
 Echo (ping) request  id=0x002a, seq=1/256, ttl=64 (no response found!)
      132 1995.231275    192.168.1.1           192.168.1.107         ICMP     132
 Destination unreachable (Host unreachable)
```

```
vm7:~ ping -c 1 192.168.8.109: error (regular)
```

Same as `vm7:~ # ping -c 1 192.168.8.108 (#vm7--ping--c-1-1921688108-error)`.

```
vm7:~ ping -c 1 192.168.9.108: error (regular)
```

Same as `vm7:~ # ping -c 1 192.168.8.108 (#vm7--ping--c-1-1921688108-error)`.

```
vm7:~ ping -c 1 192.168.9.109: error (regular)
```

Same as `vm7:~ # ping -c 1 192.168.8.108 (#vm7--ping--c-1-1921688108-error)`.

```
vm7:~ ping -I eth7 -c 1 192.168.8.108: reply (enforced)
```

Pinging an IP *outside* of a directly connected network, while enforcing the use of the specified `eth7` interface, is done instead of looking up the destination in the routing table to select an interface and source IP address. In other words, it ignores the routing table.
This causes the machine to resolve the MAC address of the machine having the IP address via ARP on the `eth7` interface and to use the primary IP address on this interface as the source IP. The destination machine has a route back to the source IP network and can deliver the reply.
Output:

```
  PING 192.168.8.108 (192.168.8.108) from 192.168.7.107 eth7: 56(84) bytes of data.
  64 bytes from 192.168.8.108: icmp_seq=1 ttl=64 time=0.431 ms

  --- 192.168.8.108 ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.431/0.431/0.431/0.000 ms
```

```
- capture:

  No.     Time           Source                Destination           Protocol Length
 Info
       55 1090.280252902 52:54:00:00:07:07     Broadcast             ARP      42
 Who has 192.168.8.108? Tell 192.168.7.107
       56 1090.280444897 52:54:00:00:07:08     52:54:00:00:07:07     ARP      42
 192.168.8.108 is at 52:54:00:00:07:08
       57 1090.280492377 192.168.7.107         192.168.8.108         ICMP     98
 Echo (ping) request  id=0x000c, seq=1/256, ttl=64 (reply in 58)
       58 1090.280553072 192.168.8.108         192.168.7.107         ICMP     98
 Echo (ping) reply    id=0x000c, seq=1/256, ttl=64 (request in 57)
       59 1095.512647002 52:54:00:00:07:08     52:54:00:00:07:07     ARP      42
 Who has 192.168.7.107? Tell 192.168.7.108
       60 1095.512893801 52:54:00:00:07:07     52:54:00:00:07:08     ARP      42
 192.168.7.107 is at 52:54:00:00:07:07

- arp-cache:

  192.168.8.108 dev eth7 lladdr 52:54:00:00:07:08 REACHABLE
  192.168.7.108 dev eth7 lladdr 52:54:00:00:07:08 REACHABLE

- strace:

  socket(AF_INET, SOCK_DGRAM, IPPROTO_ICMP) = 3
  setsockopt(3, SOL_SOCKET, SO_BINDTODEVICE, "eth7\0", 5) = 0
  sendto(3, "\10\0\2\371\377\377\0\1J\274sh\0\0\0\0j
\16\16\0\0\0\0\0\20\21\22\23\24\25\26\27"..., 64, 0, {sa_family=AF_INET, sin
  _port=htons(0), sin_addr=inet_addr("192.168.8.108")}, 16) = 64
  recvmsg(3, {msg_name={sa_family=AF_INET, sin_port=htons(0),
 sin_addr=inet_addr("192.168.8.108")}, msg_namelen=128 => 16, msg_i
  ov=[{iov_base="\0\0\n\346\0\23\0\1J\274sh\0\0\0\0j
\16\16\0\0\0\0\0\20\21\22\23\24\25\26\27"..., iov_len=192}], msg_iovlen=1, m
  sg_control=[{cmsg_len=32, cmsg_level=SOL_SOCKET, cmsg_type=SO_TIMESTAMP_OLD,
 cmsg_data={tv_sec=1752415306, tv_usec=921573}}, {
  cmsg_len=20, cmsg_level=SOL_IP, cmsg_type=IP_TTL, cmsg_data=[64]}],
 msg_controllen=56, msg_flags=0}, 0) = 64
  write(1, "64 bytes from 192.168.8.108: icm"..., 61) = 61
```

```
vm7:~ ping -I eth7 -c 1 192.168.8.109: reply (enforced)
```

Same as vm7:~ # ping -I eth7 -c 1 192.168.8.108 (#vm7--ping--i-eth7--c-1-1921688108-reply).

```
vm7:~ ping -I eth7 -c 1 192.168.9.108: reply (enforced)
```

Same as `vm7:~ # ping -I eth7 -c 1 192.168.8.108 (#vm7--ping--i-eth7--c-1-1921688108-reply)`.

```
vm7:~ # ping -I eth7 -c 1 192.168.9.109: reply (enforced)
```

Same as `vm7:~ # ping -I eth7 -c 1 192.168.8.108 (#vm7--ping--i-eth7--c-1-1921688108-reply)`.

2. vm8

   In our setup, `vm8` and `vm9` behave the same way. See `vm9` (#vm9) for more details.

3. vm9

   ```
   vm9:~ ping -c 1 192.168.7.107: reply (regular)
   ```

   Regular pinging of an IP on a directly connected network (network route) triggers a routing table lookup that selects the `enp7s0` interface and the preferred source from the route (src 192.168.7.109):

   ```
   192.168.7.0/24 dev enp7s0 proto kernel scope link src 192.168.7.109 metric 100
   ```

   Output:

   ```
     PING 192.168.7.107 (192.168.7.107) 56(84) bytes of data.
     64 bytes from 192.168.7.107: icmp_seq=1 ttl=64 time=0.567 ms

     --- 192.168.7.107 ping statistics ---
     1 packets transmitted, 1 received, 0% packet loss, time 0ms
     rtt min/avg/max/mdev = 0.567/0.567/0.567/0.000 ms

   - capture:

     No.     Time           Source                Destination           Protocol Length
    Info
           1 0.000000000    52:54:00:00:07:09     Broadcast             ARP      42
    Who has 192.168.7.107? Tell 192.168.7.109
           2 0.000356939    52:54:00:00:07:07     52:54:00:00:07:09     ARP      42
    192.168.7.107 is at 52:54:00:00:07:07
           3 0.000453833    192.168.7.109         192.168.7.107         ICMP     98
    Echo (ping) request  id=0x0006, seq=1/256, ttl=64 (reply in 4)
           4 0.000512905    192.168.7.107         192.168.7.109         ICMP     98
    Echo (ping) reply    id=0x0006, seq=1/256, ttl=64 (request in 3)
           5 23.564217607   52:54:00:00:07:07     52:54:00:00:07:09     ARP      42
    Who has 192.168.7.109? Tell 192.168.7.107
           6 23.564482951   52:54:00:00:07:09     52:54:00:00:07:07     ARP      42
    192.168.7.109 is at 52:54:00:00:07:09
   ```

```
- arp-cache:

  192.168.7.107 dev enp7s0 lladdr 52:54:00:00:07:07 REACHABLE
```

```
vm9:~ ping -c 1 192.168.7.108: reply (regular)
```

Same as vm9:~ # ping -c 1 192.168.7.107(#vm9--ping--c-1-1921687107-reply).

```
vm9:~ ping -c 1 192.168.8.108: reply (regular)
```

Regular pinging of an IP on a directly connected network (network route) triggers a routing table lookup that selects the enp8s0 interface and the preferred source from the route (src 192.168.8.109):

```
192.168.8.0/24 dev enp8s0 proto kernel scope link src 192.168.8.109 metric 101
```

Output:

```
PING 192.168.8.108 (192.168.8.108) 56(84) bytes of data.
  64 bytes from 192.168.8.108: icmp_seq=1 ttl=64 time=0.460 ms

  --- 192.168.8.108 ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.460/0.460/0.460/0.000 ms

- capture:

  No.     Time           Source                Destination           Protocol Length
 Info
        1 0.000000000    52:54:00:00:08:09     Broadcast             ARP      42
 Who has 192.168.8.108? Tell 192.168.8.109
        2 0.000155876    52:54:00:00:08:08     52:54:00:00:08:09     ARP      42
 192.168.8.108 is at 52:54:00:00:08:08
        3 0.000268681    192.168.8.109         192.168.8.108         ICMP     98
 Echo (ping) request  id=0x0008, seq=1/256, ttl=64 (reply in 4)
        4 0.000336540    192.168.8.108         192.168.8.109         ICMP     98
 Echo (ping) reply    id=0x0008, seq=1/256, ttl=64 (request in 3)
        5 5.109132370    52:54:00:00:08:08     52:54:00:00:08:09     ARP      42
 Who has 192.168.8.109? Tell 192.168.8.108
        6 5.109320257    52:54:00:00:08:09     52:54:00:00:08:08     ARP      42
 192.168.8.109 is at 52:54:00:00:08:09

- arp-cache:

  192.168.8.108 dev enp8s0 lladdr 52:54:00:00:08:08 REACHABLE
```

```
vm9:~ ping -c 1 192.168.9.108: reply (regular)
```

```
Regular pinging of an IP on a directly connected network (network route) and using
a lookup of the destination IP via routing table causing to select the enp9s0
interface and the preferred source from the route (src 192.168.9.109):
```

Output:

```
 PING 192.168.9.108 (192.168.9.108) 56(84) bytes of data.
  64 bytes from 192.168.9.108: icmp_seq=1 ttl=64 time=0.488 ms

  --- 192.168.9.108 ping statistics ---
  1 packets transmitted, 1 received, 0% packet loss, time 0ms
  rtt min/avg/max/mdev = 0.488/0.488/0.488/0.000 ms

- capture:

  No.     Time           Source                Destination           Protocol Length
 Info
        1 0.000000000    52:54:00:00:09:09     Broadcast             ARP      42
 Who has 192.168.9.108? Tell 192.168.9.109
        2 0.000212864    52:54:00:00:09:08     52:54:00:00:09:09     ARP      42
 192.168.9.108 is at 52:54:00:00:09:08
        3 0.000302044    192.168.9.109         192.168.9.108         ICMP     98
 Echo (ping) request  id=0x0009, seq=1/256, ttl=64 (reply in 4)
        4 0.000364673    192.168.9.108         192.168.9.109         ICMP     98
 Echo (ping) reply    id=0x0009, seq=1/256, ttl=64 (request in 3)
        5 5.157113256    52:54:00:00:09:08     52:54:00:00:09:09     ARP      42
 Who has 192.168.9.109? Tell 192.168.9.108
        6 5.157392878    52:54:00:00:09:09     52:54:00:00:09:08     ARP      42
 192.168.9.109 is at 52:54:00:00:09:09

- arp-cache:

  192.168.9.108 dev enp9s0 lladdr 52:54:00:00:09:08 REACHABLE
```

```
vm9:~ # ping -I enp9s0 -c 1 192.168.8.108: reply (enforced)
```

Pinging an IP address while *binding* the ping to the specified `enp9s0` interface is done instead of looking up the destination in the routing table. In other words, it ignores the routing table.
This causes the machine to resolve the MAC address of the machine having the IP address via ARP on the `enp9s0` interface and to use the primary IP address (`192.168.9.109`) as the source.

The destination machine answers with the MAC address of `enp8s0` on `enp9s0`, causing an additional MAC and IP ARP association in caches, and sends the ICMP echo reply using the same interface as in the request.
Output:

```
PING 192.168.8.108 (192.168.8.108) from 192.168.9.109 enp9s0: 56(84) bytes of data.
   64 bytes from 192.168.8.108: icmp_seq=1 ttl=64 time=0.895 ms

   --- 192.168.8.108 ping statistics ---
   1 packets transmitted, 1 received, 0% packet loss, time 0ms
   rtt min/avg/max/mdev = 0.895/0.895/0.895/0.000 ms

 - capture (net9):

   No.     Time           Source                Destination           Protocol
Length Info
        13 2139.146557619 52:54:00:00:09:09     Broadcast             ARP      42
  Who has 192.168.8.108? Tell 192.168.9.109
        14 2139.146934545 52:54:00:00:09:08     52:54:00:00:09:09     ARP      42
  192.168.8.108 is at 52:54:00:00:09:08
        15 2139.146993788 192.168.9.109         192.168.8.108         ICMP     98
  Echo (ping) request  id=0x0015, seq=1/256, ttl=64 (reply in 16)
        16 2139.147062178 192.168.8.108         192.168.9.109         ICMP     98
  Echo (ping) reply    id=0x0015, seq=1/256, ttl=64 (request in 15)
        17 2144.347165844 52:54:00:00:09:08     52:54:00:00:09:09     ARP      42
  Who has 192.168.9.109? Tell 192.168.9.108
        18 2144.347432020 52:54:00:00:09:09     52:54:00:00:09:08     ARP      42
  192.168.9.109 is at 52:54:00:00:09:09

 - arp-cache:
   (when used ping with/without interface binding multiple times)

   192.168.8.108 dev enp9s0 lladdr 52:54:00:00:09:08 REACHABLE
   192.168.8.108 dev enp8s0 lladdr 52:54:00:00:08:08 REACHABLE
```

```
vm9:~ ping -I enp8s0 -c 1 192.168.9.108: reply (enforced)
```

Same as `vm9:~ # ping -I enp9s0 -c 1 192.168.8.108 (#vm9--ping--i-enp9s0--c-1-1921688108-reply)`.
just binding the `enp8s0` interface to ping IP on `enp9s0`.

```
vm9:~ ping -I enp7s0 -c 1 192.168.9.108: error (regular)
```

Pinging an IP address bound/enforced to the specified `enp7s0` interface is done instead of looking up the destination in the routing table. In other words, it ignores the routing table.
Similar to `vm7:~ # ping -c 1 192.168.8.108 (#vm7--ping--c-1-1921688108-error)`, failing *regardless* of the interface binding because the router on the `enp7s0` interface does not have routing to the 192.168.9.0/24 network and routes the packet to the uplink.
Output (incl. unreachable error):

```
PING 192.168.9.108 (192.168.9.108) from 192.168.7.109 enp7s0: 56(84) bytes of data.
  From 192.168.1.1 icmp_seq=1 Destination Host Unreachable

  --- 192.168.9.108 ping statistics ---
  1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

- capture (net7):

  No.     Time           Source                Destination           Protocol Length
 Info
       35 468.566259392  192.168.7.109         192.168.9.108         ICMP     98
 Echo (ping) request  id=0x000c, seq=1/256, ttl=64 (no response found!)
       36 471.671403634  192.168.1.1           192.168.7.109         ICMP     126
 Destination unreachable (Host unreachable)

- arp-cache: no entry (on enp7s0)
```

```
vm9:~ # ping -I enp7s0 -c 1 192.168.8.108: error (regular)
```

Same as `vm9:~ # ping -I enp7s0 -c 1 192.168.9.108 (#vm9--ping--i-enp7s0--c-1-1921689108-error)`.

6 Multihome setup

This section shows how to configure a multihome system using kernel ARP and reverse path filtering settings to control network traffic.

6.1 Setting up multihoming

The kernel's multihome ARP filtering is controlled by the following `sysctl` settings, which are typically placed in `/etc/sysctl.d/90-network.conf` config:

PROCEDURE 4: MULTIHOME SETUP

1. Mode for announcing local IP in requests:

   ```
   net.ipv4.conf.all.arp_announce = 2
   ```

2. Mode for sending replies to ARP requests:

   ```
   net.ipv4.conf.all.arp_ignore = 1
   ```

   or `higher` filtering levels (`arp_ignore = 2` if desired) on all involved machines `vm7`, `vm8`, `vm9`. The `rp_filter = 2` filtering of IP packets according to the route table matches is already set by SUSE by default to `rp_filter = 2`. When desired, in can be “increased” to stricter filtering using `rp_filter = 1`, but please note that this may cause packet drops, for example, in an asymmetric routing setup.

3. The following ping commands illustrate the multihome behavior and difference to the default behavior:

   a. vm7

      ```
      vm7:~ ping -c 1 192.168.7.108`: reply (regular)
      ```

      Same as default behavior (#vm7--ping--c-1-1921687108-reply).

      ```
      vm7:~ ping -c 1 192.168.7.109`: reply (regular)
      ```

      Same as default behavior (#vm7--ping--c-1-1921687108-reply).

      ```
      vm7:~ # ping -c 1 192.168.8.108`: error (regular)
      ```

      Same as default behavior (#vm7--ping--c-1-1921688108-error).

      ```
      vm7:~ # ping -c 1 192.168.8.109`: error (regular)
      ```

      Same as default behavior (#vm7--ping--c-1-1921688108-error).

      ```
      vm7:~ # ping -c 1 192.168.9.108`: error (regular)
      ```

Same as default behavior (#vm7--ping--c-1-1921688108-error).

```
vm7:~ # ping -c 1 192.168.9.109`: error (regular)
```

Same as default behavior (#vm7--ping--c-1-1921688108-error).

```
vm7:~ # ping -I eth7 -c 1 192.168.8.108`: error (filtered)
```

The multihome ARP settings cause the system to ignore ARP requests due to a subnet mismatch, and as a result, the enforcement via the `-I` interface binding no longer works.
Output:

```
PING 192.168.8.108 (192.168.8.108) from 192.168.7.107 eth7: 56(84) bytes of
 data.
From 192.168.7.107 icmp_seq=1 Destination Host Unreachable

--- 192.168.8.108 ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

- capture:

No.     Time           Source                Destination           Protocol
 Length Info
     56 461.814511760  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.8.108? Tell 192.168.7.107
     57 462.819873108  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.8.108? Tell 192.168.7.107
     58 463.844043402  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.8.108? Tell 192.168.7.107
```

```
vm7:~ # ping -I eth7 -c 1 192.168.8.109`: error (filtered)
```

The multihome ARP settings cause the system to ignore ARP requests due to a subnet mismatch, and as a result, the enforcement via the `-I` interface binding no longer works.
Output:

```
PING 192.168.8.109 (192.168.8.109) from 192.168.7.107 eth7: 56(84) bytes of
 data.
From 192.168.7.107 icmp_seq=1 Destination Host Unreachable

--- 192.168.8.109 ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

- capture:
```

```
No.     Time           Source                Destination           Protocol
 Length Info
     74 623.897407366  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.8.109? Tell 192.168.7.107
     75 624.903906694  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.8.109? Tell 192.168.7.107
     76 625.928060996  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.8.109? Tell 192.168.7.107
```

```
vm7:~ ping -I eth7 -c 1 192.168.9.108: error (filtered)
```

The multihome ARP settings cause the system to ignore ARP requests due to a subnet mismatch, and as a result, the enforcement via the `-I` interface binding no longer works.
Output:

```
PING 192.168.9.108 (192.168.9.108) from 192.168.7.107 eth7: 56(84) bytes of
 data.
From 192.168.7.107 icmp_seq=1 Destination Host Unreachable

--- 192.168.9.108 ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

- capture:

No.     Time           Source                Destination           Protocol
 Length Info
     81 726.553781184  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.9.108? Tell 192.168.7.107
     82 727.562690381  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.9.108? Tell 192.168.7.107
     83 728.586537241  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.9.108? Tell 192.168.7.107
```

```
vm7:~ ping -I eth7 -c 1 192.168.9.109: error (filtered)
```

The multihome ARP settings cause the system to ignore ARP requests due to a subnet mismatch, and as a result, the enforcement via the `-I` interface binding no longer works.
Output:

```
PING 192.168.9.109 (192.168.9.109) from 192.168.7.107 eth7: 56(84) bytes of
 data.
From 192.168.7.107 icmp_seq=1 Destination Host Unreachable

--- 192.168.9.109 ping statistics ---
```

```
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

- capture:

No.     Time           Source                Destination           Protocol
 Length Info
     92 817.265392026  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.9.109? Tell 192.168.7.107
     93 818.284912731  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.9.109? Tell 192.168.7.107
     94 819.309070330  52:54:00:00:07:07     Broadcast             ARP      42
     Who has 192.168.9.109? Tell 192.168.7.107
```

b. vm8

In our setup, `vm8` and `vm9` behave the same way. See `vm9` (#vm9-1) for more details.

c. vm9

```
vm9:~ # ping -c 1 192.168.7.107: reply (regular)
```

Same as default behavior (#vm9--ping--c-1-1921687107-reply-regular).

```
vm9:~ # ping -c 1 192.168.7.108: reply (regular)
```

Same as default behavior (#vm9--ping--c-1-1921687107-reply-regular).

```
vm9:~ # ping -c 1 192.168.8.108 : reply (regular)
```

Same as default behavior (#vm9--ping--c-1-1921688108-reply-regular).

```
vm9:~ # ping -c 1 192.168.9.108 : reply (regular)
```

Same as default behavior (#vm9--ping--c-1-1921689108-reply-regular).

```
vm9:~ ping -I enp9s0 -c 1 192.168.8.108 : error (filtered)
```

The multihome ARP settings cause the system to ignore ARP requests due to a subnet mismatch, and as a result, the enforcement via the `-I` interface binding no longer works.
Output:

```
PING 192.168.8.108 (192.168.8.108) from 192.168.9.109 enp9s0: 56(84) bytes of
 data.
From 192.168.9.109 icmp_seq=1 Destination Host Unreachable

--- 192.168.8.108 ping statistics ---
```

```
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

- capture:

No.     Time           Source                Destination           Protocol
 Length Info
      1 0.000000000    52:54:00:00:09:09     Broadcast             ARP      42
     Who has 192.168.8.108? Tell 192.168.9.109
      2 1.053440783    52:54:00:00:09:09     Broadcast             ARP      42
     Who has 192.168.8.108? Tell 192.168.9.109
      3 2.077406483    52:54:00:00:09:09     Broadcast             ARP      42
     Who has 192.168.8.108? Tell 192.168.9.109
```

```
vm9:~ ping -I enp8s0 -c 1 192.168.9.108 : error (filtered)
```

The multihome ARP settings cause the system to ignore ARP requests due to a subnet mismatch, and as a result, the enforcement via the `-I` interface binding no longer works.
Output:

```
PING 192.168.9.108 (192.168.9.108) from 192.168.8.109 enp8s0: 56(84) bytes of
 data.
From 192.168.8.109 icmp_seq=1 Destination Host Unreachable

--- 192.168.9.108 ping statistics ---
1 packets transmitted, 0 received, +1 errors, 100% packet loss, time 0ms

- capture:

No.     Time           Source                Destination           Protocol
 Length Info
      1 0.000000000    52:54:00:00:08:09     Broadcast             ARP      42
     Who has 192.168.9.108? Tell 192.168.8.109
      2 1.030989481    52:54:00:00:08:09     Broadcast             ARP      42
     Who has 192.168.9.108? Tell 192.168.8.109
      3 2.054972421    52:54:00:00:08:09     Broadcast             ARP      42
     Who has 192.168.9.108? Tell 192.168.8.109
```

```
vm9:~ ping -I enp7s0 -c 1 192.168.9.108: error (regular)
```

Same as default behavior (#vm9--ping--c-1-1921689108-reply-regular).

```
vm9:~ ping -I enp7s0 -c 1 192.168.8.108: error (regular)
```

Same as default behavior (#vm9--ping--c-1-1921689108-reply-regular).

6.2 Policy routing

By default, all unicast routing entries are placed in the `main` table (and local/broadcast routes `local` table) and the lookups use the destination only.

The direct routes (without a gateway) are created automatically by the kernel while adding the address to the interface, e.g. (on `vm7`). **`ip addr add 192.168.7.107/24 dev eth7`** causes to create the following routes.

```
ip route show table main dev eth7
  192.168.7.0/24 dev eth7 proto kernel scope link src 192.168.7.107
```

```
ip route show table local dev eth7

    local 192.168.7.107 proto kernel scope host src 192.168.7.107
    broadcast 192.168.7.255 proto kernel scope link src 192.168.7.107
```

Using a `/32` IP address prefix length **`ip addr add 192.168.7.107/32`** causes the kernel to avoid adding routes automatically in favor of own routes.

In complex setups involving asynchronous routing or the use of multiple interfaces in the same subnet, `arp_filter=1` may not work properly due to the filtering, especially in the strict `rp_filter=1` mode.

Policy routing permits adding a rule that matches the source IP and then using a different routing table (for example, one for each interface) that has different direct and gateway routes than the standard `main` table, such as its own unique default gateway.

For more information, see also LARTC HOWTO (https://lartc.org/lartc.html#AEN267), the `man 5 ifrule`, `man 8 ip-rule`, `man 5 ifroute`, `man 8 ip-route` man pages and the `/etc/iproute2/rt_tables` file, which is used to define names for custom tables.

7 Legal Notice

Copyright© 2006– 2026 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled “GNU Free Documentation License”.

For SUSE trademarks, see https://www.suse.com/company/legal/. All other third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors, nor the translators shall be held liable for possible errors or the consequences thereof.

A GNU Free Documentation License

Copyright (C) 2000, 2001, 2002 Free Software Foundation, Inc. 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA. Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or non-commercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or non-commercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download

using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.

C. State on the Title page the name of the publisher of the Modified Version, as the publisher.

D. Preserve all the copyright notices of the Document.

E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.

F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.

H. Include an unaltered copy of this License.

I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.

J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.

L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.

M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.

N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.

O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties--for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See https://www.gnu.org/copyleft/.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been

published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

```
Copyright (c) YEAR YOUR NAME.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included in the section entitled “GNU
Free Documentation License”.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with...Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the
Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.