

Linux Network Naming Schemes

Planning / Implementation

In the 1990s, network names were assigned by kernel, typically as eth0, eth1, and so on. However, as the boot process became increasingly parallel, these network names were no longer fixed or predictable. To address this issue, starting with version 197, systemd/udev introduced an improved naming scheme based on MAC addresses and the physical locations of network interfaces, ensuring consistent and stable interface names across reboots.

This paper explains various network naming schemes and provides guidance on configuring them to suit your preferences on a ThinkSystem server running Red Hat Enterprise Linux (RHEL) 9.2, SLES 15 SP4, or Ubuntu Server 22.04.

Traditional Naming Scheme

When the Operating System (OS) boots, network interfaces are assigned by kernel with format `eth[Index]`. The `[Index]` is the order number of network devices that match the order of the loaded driver modules. As the detection is in parallel, the assigned names may not be consistent after reboot, especially on a server with multiple network interfaces. This will cause the network to be unavailable due to firewall rules and other limitations. Most of the Linux distributions already change the default naming scheme to “Predictable naming” scheme, such as RHEL and Ubuntu, as described in the next section. However, for the users who prefer the ethX format or only have one network port, they still can use this scheme by passing the parameter `net.ifnames=0` on the kernel command line.

Predictable Naming Scheme

This section introduces the benefits and workflow of predictable naming schemes.

Benefits

The users can get the following benefits from this new scheme compared to the traditional naming scheme:

1. Consistent network port names across reboots.
2. Compatible with all the Linux distributions with systemd/udev.
3. The names are predictable as it maps to a unique attribute, such as MAC addresses, physical slots of the hardware, and others.
4. Easy to customize the names for the appropriate ports.
5. Easy to switch to another naming scheme.

Workflow

The udev manager defines the network name via 99-default.link files located in the `/usr/lib/systemd/network/` and the local administration directory `/etc/systemd/network/` (Files in `/etc/` have the highest priority), once a higher priority rule is used, the lower priority rules will not run.

```
rename:~ # grep NamePolicy /usr/lib/systemd/network/99-default.link
NamePolicy=keep kernel database onboard slot path
```

Figure 1. The name policy order of 99-default.link

The table below introduces the definition of each name policy.

Table 1. Name policy definition

| Policy | Property/Configuration File | Example | Description |
|-------------------------------|--|-----------------|---|
| Customized udev rules | /etc/systemd/network/70-*.link | PXE | Defines a custom name for the specified device. |
| The biosdevname rules | /usr/lib/udev/rules.d/71-biosdevname.rules | em1 | A predictable naming rule created by Dell. |
| kernel | N/A | lo | If kernel detects the device name is predictable, it will be skipped. |
| database | ID_NET_NAME_FROM_DATABASE | idrac | Only have an entry. |
| onboard | ID_NET_NAME_ONBOARD | eno0 | Based on index number. |
| slot | ID_NET_NAME_SLOT | ens9f0 | Based on the PCI Express hotplug network devices with available slot index number. |
| path | ID_NET_NAME_PATH | enp33s0f0 | Based on the physical/geographical location of the hardware. |
| MAC address | ID_NET_NAME_MAC | enxAABBCCDDEEFF | It does not used by default. |
| Traditional kernel-based name | N/A | eth0 | If none of the above polices can be met, it will apply the unpredictable naming scheme. |

The predictable naming structure is based on the interface types (represented by a two-character prefix) and supported attributes. Table 2 outlines the components that make up a fixed name.

Table 2. Predictable Naming Scheme

| Interface Type | Prefix | | Name Policy | Format (abbreviation) |
|----------------|--------|---|-------------|---------------------------|
| ethernet | en | + | Onboard | o (Lowercase O, not zero) |
| wlan | wl | | Slot | s [f] [d] |
| wwan | ww | | Path | [P]ps [f] [d] |
| | | | All | x |

Setting up the network port names in Linux

In this section, we provide a detailed guide on how to set the network name to your desired preference.

Switch to traditional naming scheme

If you want to use traditional naming scheme, all the name-assignment mechanisms need to be removed.

1. Create an empty 99-default.link file in the /etc/systemd/network/ and regenerate the initrd RAM disk file.

```
# mkdir -p /etc/systemd/network/  
# echo > /etc/systemd/network/99-default.link
```

For Red Hat and SLES:

```
# dracut -f
```

For Ubuntu:

```
# sudo update-initramfs -u -k all
```

2. Remove the above customized configuration files. For example:

```
# rm -f /usr/lib/udev/rules.d/71-biosdevname.rules /etc/systemd/network/70-persistent-net.rules
```

3. Remove the naming content in the relevant configuration files. For example, delete the content within the red rectangle in the 01-net.yml file and apply the new configuration in Ubuntu.

```
root@len:~# cat /etc/netplan/01-net.yaml  
# This file is generated from information provided by the datasource. Changes  
# to it will not persist across an instance reboot. To disable cloud-init's  
# network configuration capabilities, write a file  
# /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg with the following:  
# network: {config: disabled}  
network:  
  renderer: networkd  
  ethernets:  
    enp90s0f3np3:  
      dhcp4: true  
    my_lan:  
      dhcp4: true  
      match:  
        macaddress: a0:36:9f:34:89:39  
        set-name: my_lan  
  version: 2  
root@len:~#  
root@len:~# netplan apply
```

Figure 2. Customized netplan file

4. All the network ports will use the traditional naming scheme after reboot.

```
root@len:~# ls /sys/class/net/  
eth0 eth1 eth2 eth3 eth4 eth5 lo  
root@len:~#
```

Figure 3. Traditional network name list

Switch to predictable naming schemes

Most Linux distributions have adopted the predictable naming scheme as the default, but some, like SUSE, still use the traditional naming scheme by default. You can simply switch to the predictable naming scheme by adding boot parameter `net.ifnames=1`.

The following is an example of SLES 15 SP4:

1. Edit the `/etc/default/grub` file and regenerate the `grub.cfg` file.

```
# sed -i "/GRUB_CMDLINE_LINUX=/s/\\"net.ifnames=1 \"/ /etc/default/grub
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

2. The name of network ports should be predictable after reboot.

```
slesRename:~ # ls /sys/class/net/
br0  eno1  eno2  ens6f0  ens6f1  ens6f2  ens6f3  lo
slesRename:~ #
```

Figure 4. Predictable name list

Customized naming schemes

For a server with multiple network ports, the users may assign them to different tasks or prevent unexpected name changes by creating a custom naming scheme.

In this section:

- [Customize the name by switching the udev policy order](#)
- [Customize the name by using udev .rules files](#)
- [Customize the name by using systemd .link files](#)
- [Customize the name by using netplan files \(Ubuntu only\)](#)

Customize the name by switching the udev policy order

For a server with multiple network cards, the port names of the different network cards may flow different udev policy and have different format by default. Some names with slot policy, the other names with path policy. The users can adjust the udev policy order to have the uniform name format.

1. Copy the system configuration file to the higher priority directory.

```
# cp /usr/lib/systemd/network/99-default.link /etc/systemd/network/
```

2. Modify the 99-default.link. For example, set policy path priority over slot.

```
rename:~ # sed -i "/NamePolicy/s/\\(slot\\) \\(path\\)/\\2 \\1/" /etc/systemd/network/99-default.link
rename:~ # cat /etc/systemd/network/99-default.link |grep -2 NamePolicy

[Link]
NamePolicy=keep kernel database onboard path slot
AlternativeNamesPolicy=database onboard slot path
MACAddressPolicy=persistent
rename:~ #
```

Figure 5. Customized 99-default.link

3. Regenerate the initrd RAM disk file.
For Red Hat and SLES:

```
# dracut -f
```

For Ubuntu:

```
# sudo update-initramfs -u -k all
```

4. Reboot the OS.
5. The name scheme replaces slot to path.

```
rename:~ # ls /sys/class/net/
br0  eno1  eno2  enp13s0f0  enp13s0f1  enp13s0f2  enp13s0f3
rename:~ #
```

Figure 6. Network name list with policy "path"

Customize the name by using udev .rules files

The users can assign a fixed name for any interface via /etc/udev/rules.d/70-persistent-net.rules. The example below will use MAC address and bus info to name two different network ports.

1. Identify the device MAC address and bus info.

```
rename:~ # udevadm info -a -p /sys/class/net/enp13s0f1 |grep -Ew "type|dev_id|address"
ATTR{address}=="a0:36:9f:34:89:39"
ATTR{dev_id}=="0x0"
ATTR{type}=="1"
rename:~ #
```

Figure 7. Network port attribute (type, MAC address, dev_id)

```
rename:~ # udevadm info /sys/class/net/enp13s0f2 |grep -w ID_PATH
E: ID_PATH=pci-0000:0d:00.2
rename:~ #
```

Figure 8. The path of network port

2. Edit the file /etc/udev/rules.d/70-persistent-net.rules.

```
rename:~ # cat /etc/udev/rules.d/70-persistent-net.rules
#Assign the fixed name via bus address
SUBSYSTEM=="net", ACTION=="add", KERNELS=="0000:0d:00.2", NAME="PXE"
#Assign the fixed name via MAC address
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?* ", ATTR{address}=="a0:36:9f:34:89:39", ATTR{dev_id}=="0x0"
, ATTR{type}=="1", KERNEL=="eth*", NAME="my_lan1"
rename:~ #
```

Figure 9. Customized 70-persistent-net.rules

3. Regenerate the initrd RAM disk file
4. Reboot the OS.
5. Verify the new network name.

```
rename:~ # ip addr show my_lan1
5: my_lan1: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
    link/ether a0:36:9f:34:89:39 brd ff:ff:ff:ff:ff:ff
    altname ens6f1
    altname enp13s0f1
rename:~ #
```

Figure 10. Customized network name with udev rules and MAC address

```
rename:~ # ethtool -i PXE
driver: igb
version: 5.14.21-150400.22-default
firmware-version: 1.63, 0x80000a58
expansion-rom-version:
bus-info: 0000:0d:00.2
supports-statistics: yes
```

Figure 11. Customized network name with udev rules and path

Customize the name by using systemd .link files

In addition to the new default naming scheme, the users can customize it to fit their needs by using the `/etc/systemd/network/70-*.link` file.

1. Identify the device MAC address and bus info (Same as the steps in [Customize the name by .using udev .rules files](#) section).
2. Create a systemd configuration file for two network ports.

```
rename:~ # cat /etc/systemd/network/70-mac.link
[Match]
MACAddress=a0:36:9f:34:89:39

[Link]
Name=link_mylan
MACAddressPolicy=persistent
rename:~ #
```

Figure 12. systemd 70-mac.link

```
rename:~ # cat /etc/systemd/network/70-path.link
[Match]
Path=pci-0000:0d:00.2

[Link]
Name=link_PXE
rename:~ #
```

Figure 13. systemd 70-path.link

3. Regenerate the initrd RAM disk file
4. Reboot the OS.
5. Verify the new network name.

```
rename:~ # ip addr show link_mylan
5: link_mylan: <BROADCAST,MULTICAST> mtu 1500 qdisc noop state DOWN group default qlen 1000
    link/ether a0:36:9f:34:89:39 brd ff:ff:ff:ff:ff:ff
rename:~ #
```

Figure 14. Customized network name with systemd and MAC address

```
rename:~ # udevadm info /sys/class/net/link_PXE | grep -w ID_PATH
E: ID_PATH=pci-0000:0d:00.2
rename:~ #
```

Figure 15. Customized network name with systemd and path

Customize the name by using netplan files (Ubuntu only)

In Ubuntu, the software netplan also can set a fixed name for a specified MAC address.

1. Create the YAML format file for a specified MAC address and apply it.

```
root@rename:~# cat /etc/netplan/01-mac.yaml
# This file is generated from information provided by the datasource. Changes
# to it will not persist across an instance reboot. To disable cloud-init's
# network configuration capabilities, write a file
# /etc/cloud/cloud.cfg.d/99-disable-network-config.cfg with the following:
# network: {config: disabled}
network:
  renderer: networkd
  ethernets:
    my_lan_mac:
      dhcp4: true
      match:
        macaddress: a0:36:9f:34:89:39
      set-name: my_lan_mac
  version: 2
root@rename:~#
root@rename:~# netplan apply
```

Figure 16. Customized netplan file with MAC address

2. Verify the new network name.

```
root@rename:~# ip addr show my_lan_mac
5: my_lan_mac: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN group default qlen 1000
    link/ether a0:36:9f:34:89:39 brd ff:ff:ff:ff:ff:ff
    altname ens6f1
root@rename:~#
```

Figure 17. Customized network name with netplan and MAC address

Limitations

If users want to change the default naming scheme, they should also be mindful of the following limitations.

1. Be careful to use the MAC address to map the network name, as the systemd file may replace the MAC address from persistent to random via `MACAddressPolicy=random`, resulting in network's unavailability.
2. The custom network name must differ from the predefined naming schemes to avoid conflicts.

References

For more information, see these resources:

- Enabling predictable naming scheme for network interfaces
<https://documentation.suse.com/smart/network/html/network-interface-predictable-naming/index.html>
- `cmdline(7)` — Linux manual page
<https://man7.org/linux/man-pages/man7/dracut.cmdline.7.html>
- Configuring and managing networking on RHEL 9
https://docs.redhat.com/en/documentation/red_hat_enterprise_linux/9/html-single/configuring_and_managing_networking/index#assigning-alternative-names-to-a-network-interface-by-using-systemd-link-files_consistent-network-interface-device-naming
- Configuring networks on Ubuntu
<https://ubuntu.com/server/docs/configuring-networks>

Author

Song Shang is a Linux Engineer in Lenovo Infrastructure Solutions Group, based in TianJin, China.

Thanks to the following people for their contributions to this project:

- David Watts, Lenovo Press
- Adrian Huang, Lenovo Linux Engineer
- Gary Cudak, Lenovo Lead Architect

Related product families

Product families related to this document are the following:

- [Red Hat Enterprise Linux](#)
- [SUSE Linux Enterprise Server](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2025. All rights reserved.

This document, LP2175, was created or updated on March 5, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2175>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2175>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ThinkSystem®

The following terms are trademarks of other companies:

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

Other company, product, or service names may be trademarks or service marks of others.