



OpenInfra & Cloud Native
Days Korea

Linux Kernel Networking Stack

진입 장벽 허물기

강주희
claudiajkang@gmail.com



강주희 (Juhee Kang, claudiajkang@gmail.com)

- 오픈 소스 파트 프런티어 개발자
- 프로젝트: Linux Kernel Networking Stack, Kubernetes i18n
- 주력 분야 : 클라우드, 네트워크, 백엔드
- 강연 활동 : 연세대, 송실대 등 5개 대학 Git/GitHub 교육 진행

1

Linux Kernel
Networking Stack

2

Kernel 생태계

3

Kernel 기여하기

한번쯤은?


리눅스 커널 개발을 해볼까?

Linux Kernel 컨트리뷰션?



어떻게 시작하지?

다른 오픈소스처럼 컨트리뷰션 하면 될까?




리눅스 커널도 오픈소스이니까 다른 오픈 소스처럼 github에서 하겠지?

github으로 관리되는 오픈소스에 컨트리뷰션을 한다면?



오픈소스 원본 저장소

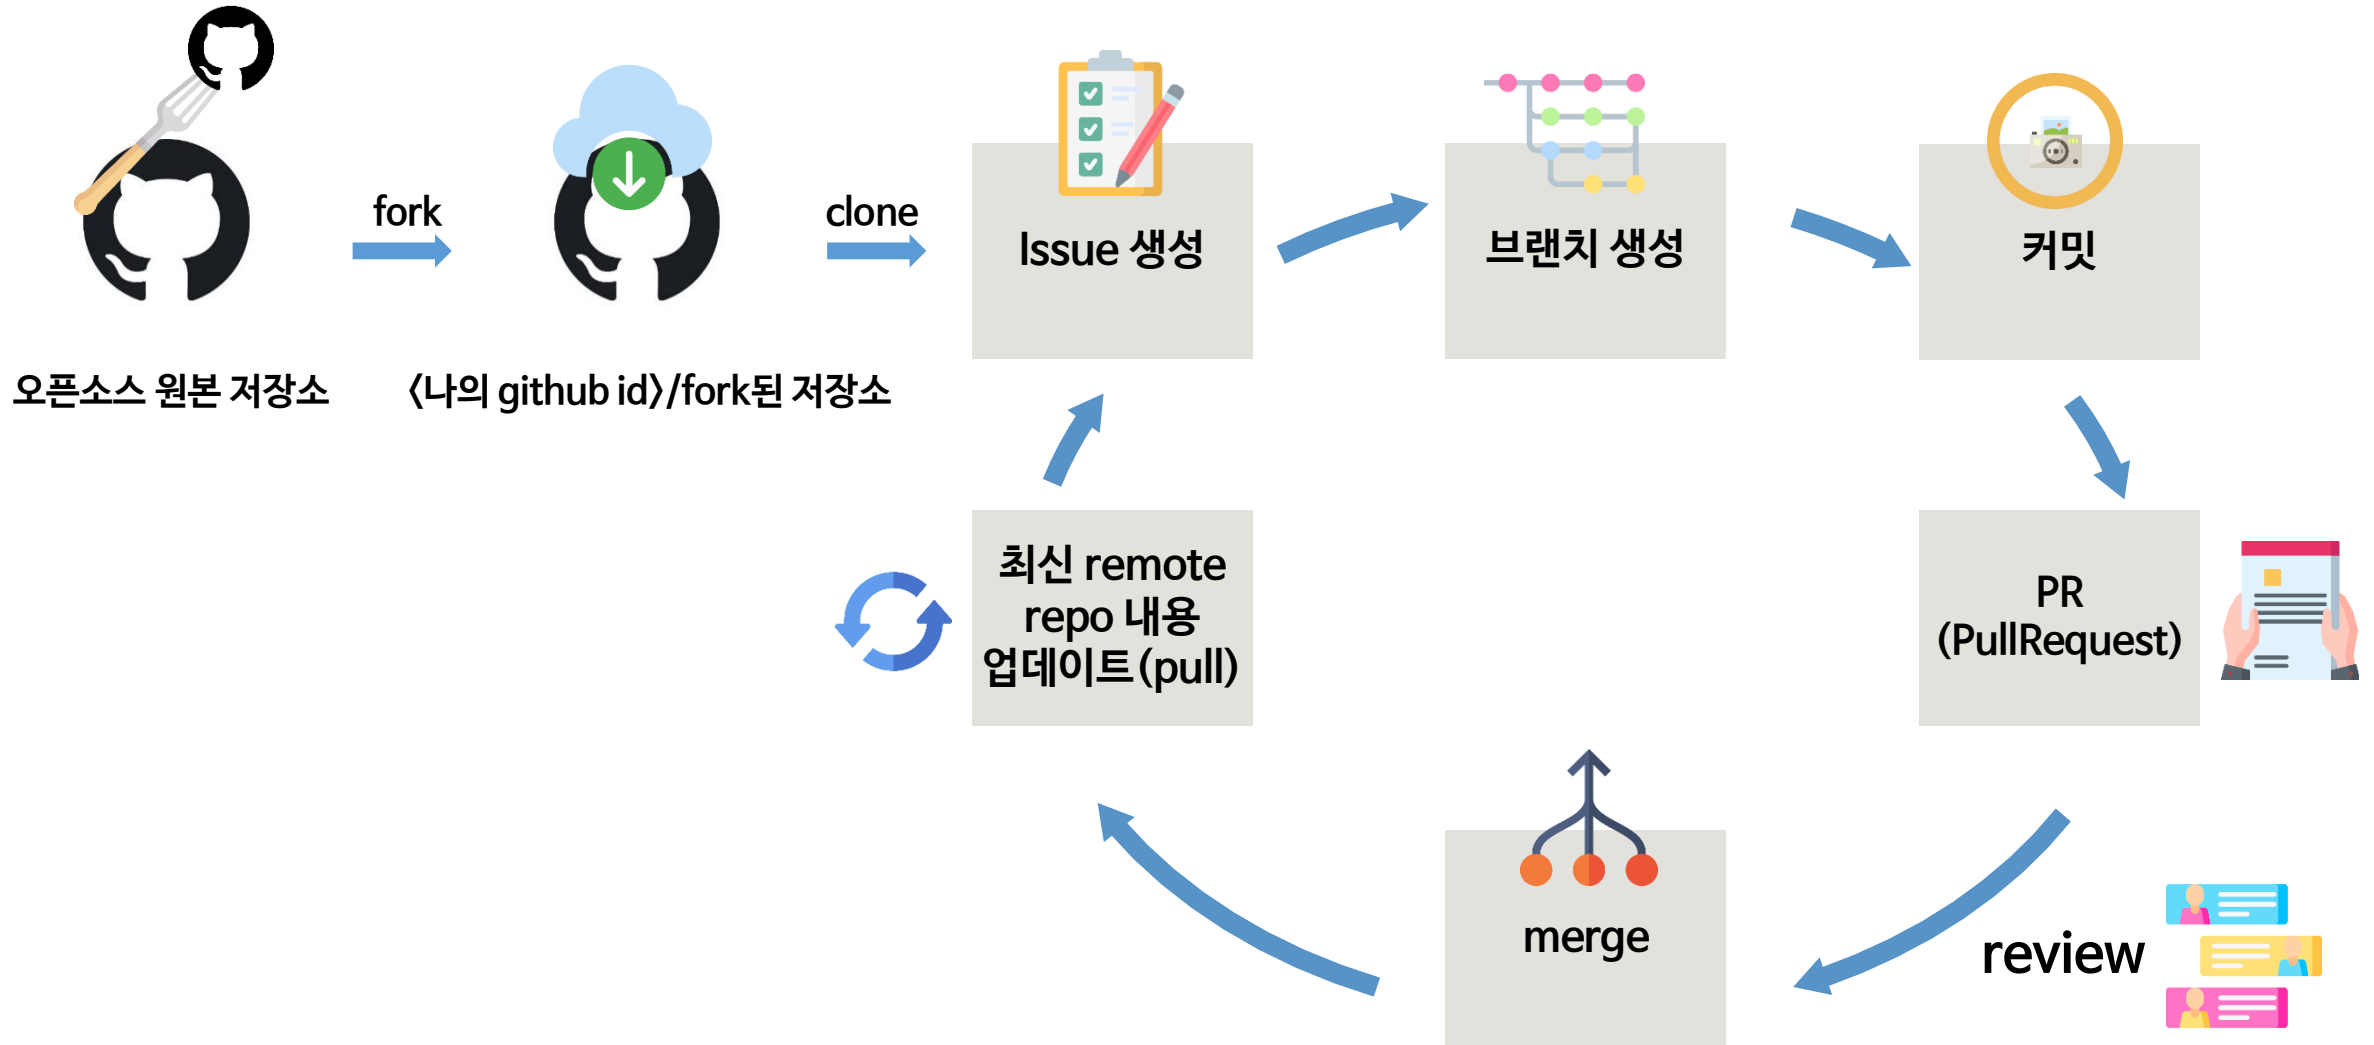
github으로 관리되는 오픈소스에 컨트리뷰션을 한다면?



오픈소스 원본 저장소

〈나의 github id〉/fork된 저장소

github으로 관리되는 오픈소스에 컨트리뷰션을 한다면?



타 오픈소스처럼 github에서 컨트리뷰션 하면 될까?

torvalds / linux Public

Watch 8.1k Fork 45.1k Starred 140k


Code Pull requests 312 Actions Projects Security Insights

master 1 branch 763 tags

Go to file

Add file

Code

 torvalds Merge tag 'for-6.1-rc1-tag' of git://git.kernel.org/pub/scm/linux/ker... aae703b 2 days ago 1,135,998 commits		
Documentation	Merge tag 'random-6.1-rc1-for-linus' of git://git.kernel.org/pub/scm/...	4 days ago
LICENSES	LICENSES/LGPL-2.1: Add LGPL-2.1-or-later as valid identifiers	10 months ago
arch	Merge tag 'random-6.1-rc1-for-linus' of git://git.kernel.org/pub/scm/...	4 days ago
block	Merge tag 'random-6.1-rc1-for-linus' of git://git.kernel.org/pub/scm/...	4 days ago
certs	certs: make system keyring depend on built-in x509 parser	27 days ago
crypto	treewide: use get_random_bytes() when possible	9 days ago
drivers	Merge tag 'v6.1-p2' of git://git.kernel.org/pub/scm/linux/kernel/git/...	3 days ago
fs	Merge tag 'for-6.1-rc1-tag' of git://git.kernel.org/pub/scm/linux/ker...	2 days ago
include	Merge tag 'cgroup-for-6.1-rc1-fixes' of git://git.kernel.org/pub/scm/...	3 days ago
init	Merge tag 'mm-nonmm-stable-2022-10-11' of git://git.kernel.org/pub/s...	8 days ago
io_uring	io_uring/rw: ensure kiocb_end_write() is always called	8 days ago
ipc	Merge tag 'mm-nonmm-stable-2022-10-11' of git://git.kernel.org/pub/s...	8 days ago

About

Linux kernel source tree

Readme

View license

140k stars

8.1k watching

45.1k forks

Releases

763 tags

Packages

No packages published

Contributors 5,000+



Github 내 kernel 소스코드 관리하는 repository로 이동

타 오픈소스처럼 github에서 컨트리뷰션 하면 될까?

torvalds / linux Public

Watch 8.1k

Fork 45.1k

Starred 140k

Code Pull requests 312 Actions Projects Security Insights

Filters is:pr is:open

Labels 0

Milestones 0

New pull request

312 Open ✓ 449 Closed

Author Label Projects Milestones Reviews Assignee Sort

- Fix for packets being rejected in the ring buffer used by the xHCI co...
#805 opened on 2 Mar 2020 by TarAldarion 1
- Update .clang-format
#803 opened on 29 Feb 2020 by farley-melo 1
- Remove Intel unavailable hardware from SDIO IDs
#802 opened on 25 Feb 2020 by laserjobs 1
- Improve readability of socket err
#801 opened on 18 Feb 2020 by sparshrestha 2
- fix "unsigned" warn when running independently
#800 opened on 18 Feb 2020 by daTokenizer 1
- Add support for USB modems using Conexant CX06836
#798 opened on 12 Feb 2020 by stevenvw1973 1
- Integration/soundwire latest
#797 opened on 11 Feb 2020 by yunbotsai 1

Pull Request 존재 확인

타 오픈소스처럼 github에서 컨트리뷰션 하면 될까?

torvalds / linux Public

Watch 8.1k

Fork 45.1k

Starred 140k

Code Pull requests 312 Actions Projects Security Insights

Fix for packets being rejected in the ring buffer used by the xHCI co... #805

Code

Open

TarAldarion wants to merge 1 commit into torvalds:master from TarAldarion:MRU_FIXES

Conversation 1

Commits 1

Checks 0

Files changed 1

+7 -0



TarAldarion commented on 2 Mar 2020

...ntroller. When a packet larger than MTU arrives in Linux from the modem, it is discarded with -EOVERFLOW error (Babble error). This is seen on USB3.0 and USB2.0 busses. This is essentially because the MRU (Max Receive Size) is not a separate entity to the MTU (Max Transmit Size) and the received packets can be larger than those transmitted. Following the babble error there were an endless supply of zero-length URBs which are rejected with -EPROTO (increasing the rx input error counter each time). This is only seen on USB3.0. These continue to come ad infinitum until the modem is shutdown. There appears to be a bug in the core USB handling code in Linux that doesn't deal well with network MTUs smaller than 1500 bytes. By default the dev->hard_mtu (the real MTU) is in lockstep with dev->rx_urb_size (essentially an MRU), and it's the latter that is causing trouble. This has nothing to do with the modems; the issue can be reproduced by getting a USB-Ethernet dongle, setting the MTU to 1430, and pinging with size greater than 1406.

Fix for packets being rejected in the ring buffer used by the xHCI co... 637f847



KernelPRBot commented on 2 Mar 2020

Hi @TarAldarion

KernelPRBot 댓글 확인

Reviewers

No reviews

Assignees

No one assigned

Labels

None yet

Projects

None yet

Milestone

No milestone

타 오픈소스처럼 컨트리뷰션 하면 될까?



Fix for packets being rejected in the ring buffer used by the xHCI co... #805
TarAldarion wants to merge 1 commit into `torvalds:master` from `TarAldarion:MRU_FIXES`



KernelPRBot commented on 2 Mar 2020

Hi @TarAldarion!

Thanks for your contribution to the Linux kernel!

Linux kernel development happens on mailing lists, rather than on GitHub - this GitHub repository is a read-only mirror that isn't used for accepting contributions. So that your change can become part of Linux, please email it to `linux@kernel.org` as a patch.

Sending patches isn't quite as simple as sending a pull request, but fortunately it is a well documented process.

Here's what to do:

- Format your contribution according to kernel requirements
- Decide who to send your contribution to
- Set up your system to send your contribution as an email
- Send your contribution and wait for feedback

How do I format my contribution?

The Linux kernel community is notoriously picky about how contributions are formatted and sent. Fortunately, they have documented their expectations.

Firstly, all contributions need to be formatted as patches. A patch is a plain text document showing the change you want to make to the code, and documenting why it is a good idea.

You can create patches with `git format-patch`.

Secondly, patches need 'commit messages', which is the human-friendly documentation explaining what the change is and why it's necessary.

need to comply with.

Github 은 read-only mirror

Milestone

No milestone

Notifications

Customize

Subscribe

You're not receiving notifications from this thread.

2 participants



타 오픈소스처럼 컨트리뷰션 하면 될까?



Fix for packets being rejected in the ring buffer used by the xHCI co... #805
TarAldarion wants to merge 1 commit into `torvalds:master` from `TarAldarion:MRU_FIXES`



KernelPRBot commented on 2 Mar 2020

Hi @TarAldarion!

Thanks for your contribution to the Linux kernel!

Linux kernel development happens on mailing lists, rather than on GitHub - this GitHub repository is a read-only mirror that isn't used for accepting contributions. So that your change can become part of Linux, please email it to us as a patch.

Sending patches isn't quite as simple as sending a pull request, but fortunately it is a well documented process.

Here's what to do:

- Format your contribution according to kernel requirements
- Decide who to send your contribution to
- Set up your system to send your contribution as an email
- Send your contribution and wait for feedback

How do I format my contribution?

The Linux kernel community is notoriously picky about how contributions are formatted and sent. Fortunately, they have documented their expectations.

Firstly, all contributions need to be formatted as patches. A patch is a plain text document showing the change you want to make to the code, and documenting why it is a good idea.

You can create patches with `git format-patch`.

Secondly, patches need 'commit messages', which is the human-friendly documentation explaining what the change is and why it's necessary.

need to comply with.

Milestone

No milestone

Notifications

Customize

Subscribe

You're not receiving notifications from this thread.

2 participants



KernelPRBot이 안내하는 Contribution 방법

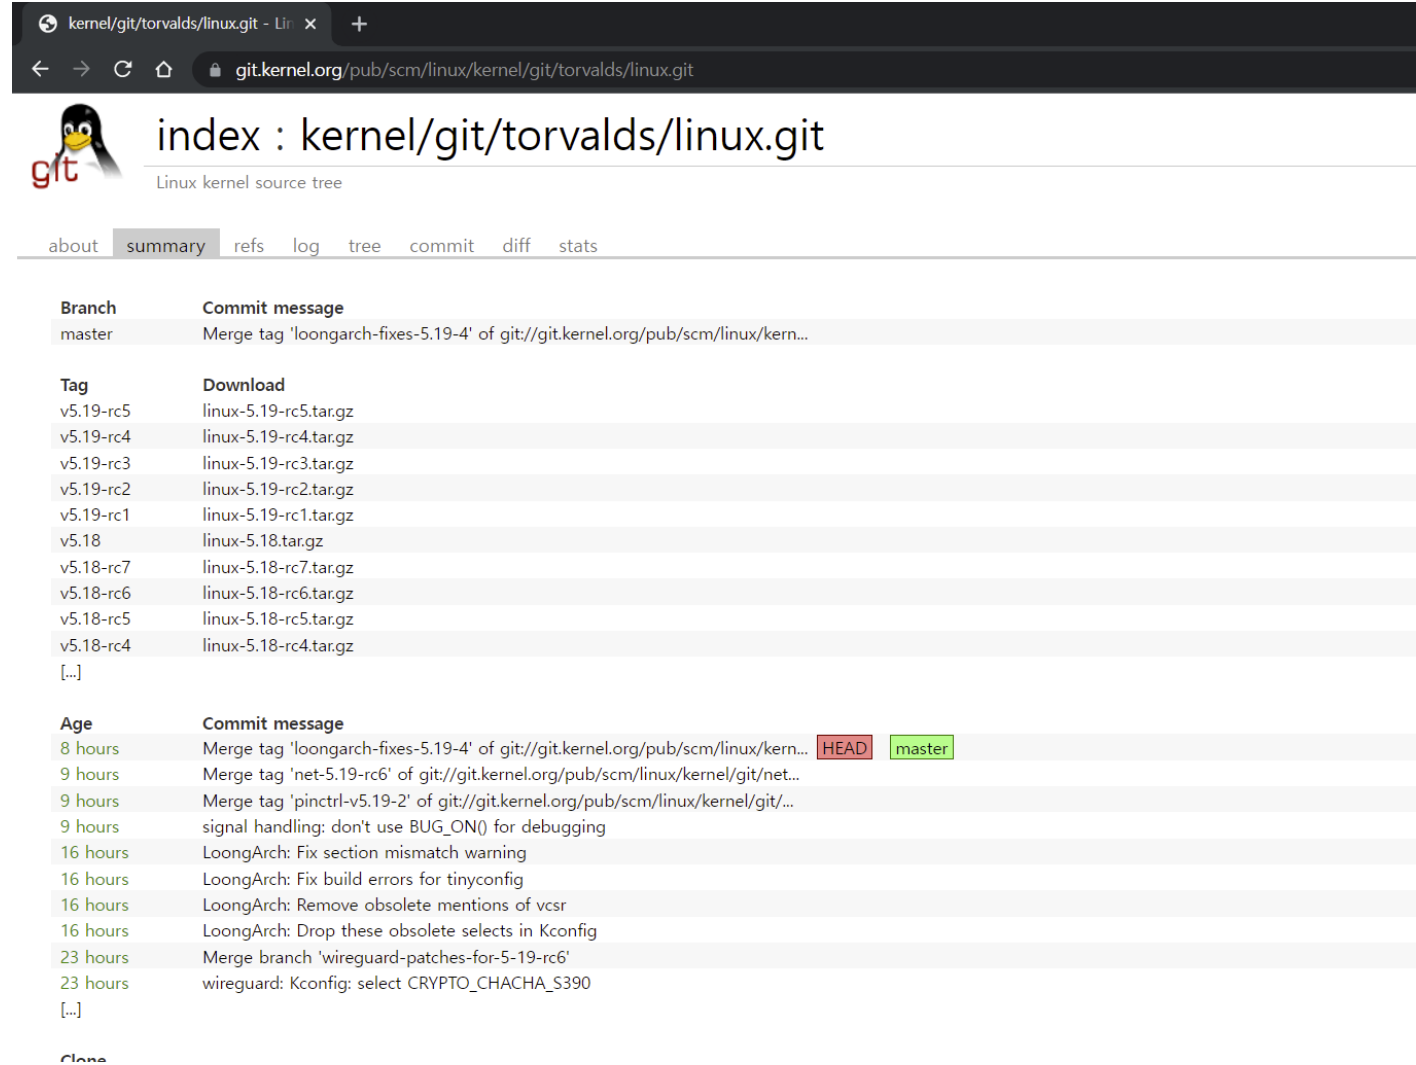
Linux Kernel 컨트리뷰션?



Linux Kernel은 GitHub을 사용하지 않는
오픈 소스 프로젝트

- Linux Kernel

(<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git>)



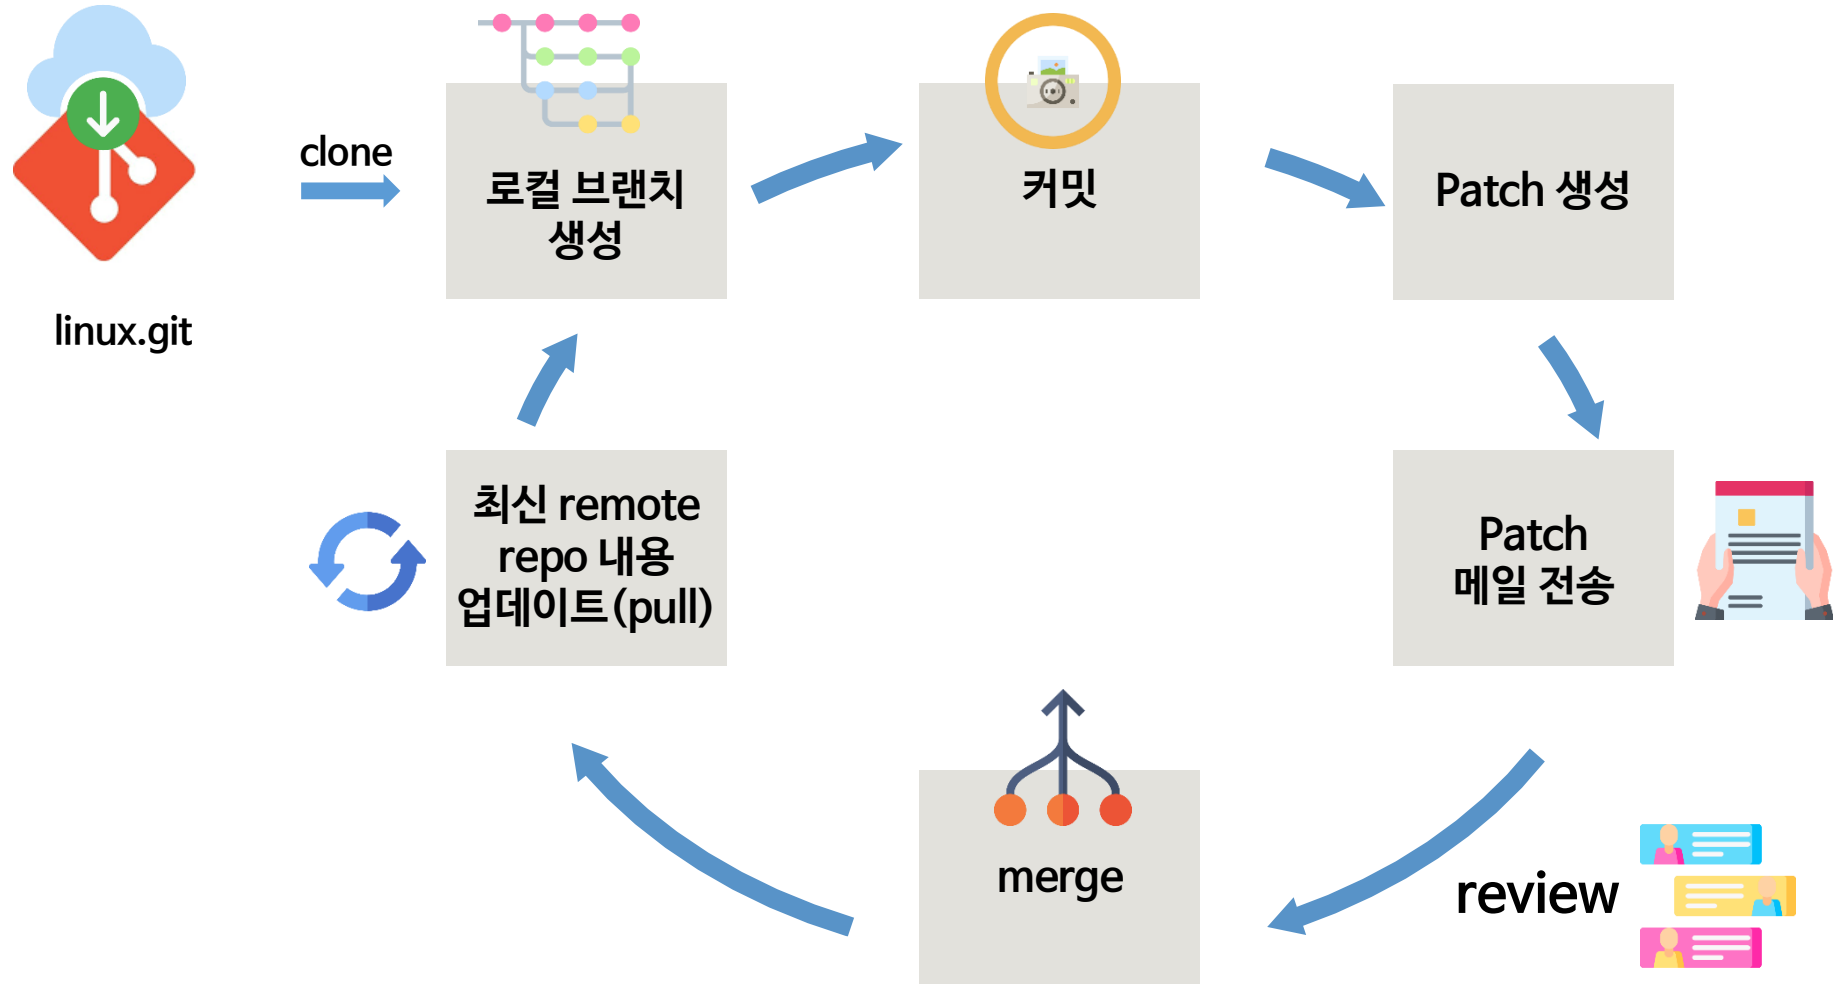
The screenshot shows the web interface for the Linux kernel source tree repository. The browser address bar displays the URL `git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git`. The page title is "index : kernel/git/torvalds/linux.git" with the subtitle "Linux kernel source tree". Navigation tabs include "about", "summary", "refs", "log", "tree", "commit", "diff", and "stats".

Branch	Commit message
master	Merge tag 'loongarch-fixes-5.19-4' of git://git.kernel.org/pub/scm/linux/kern...

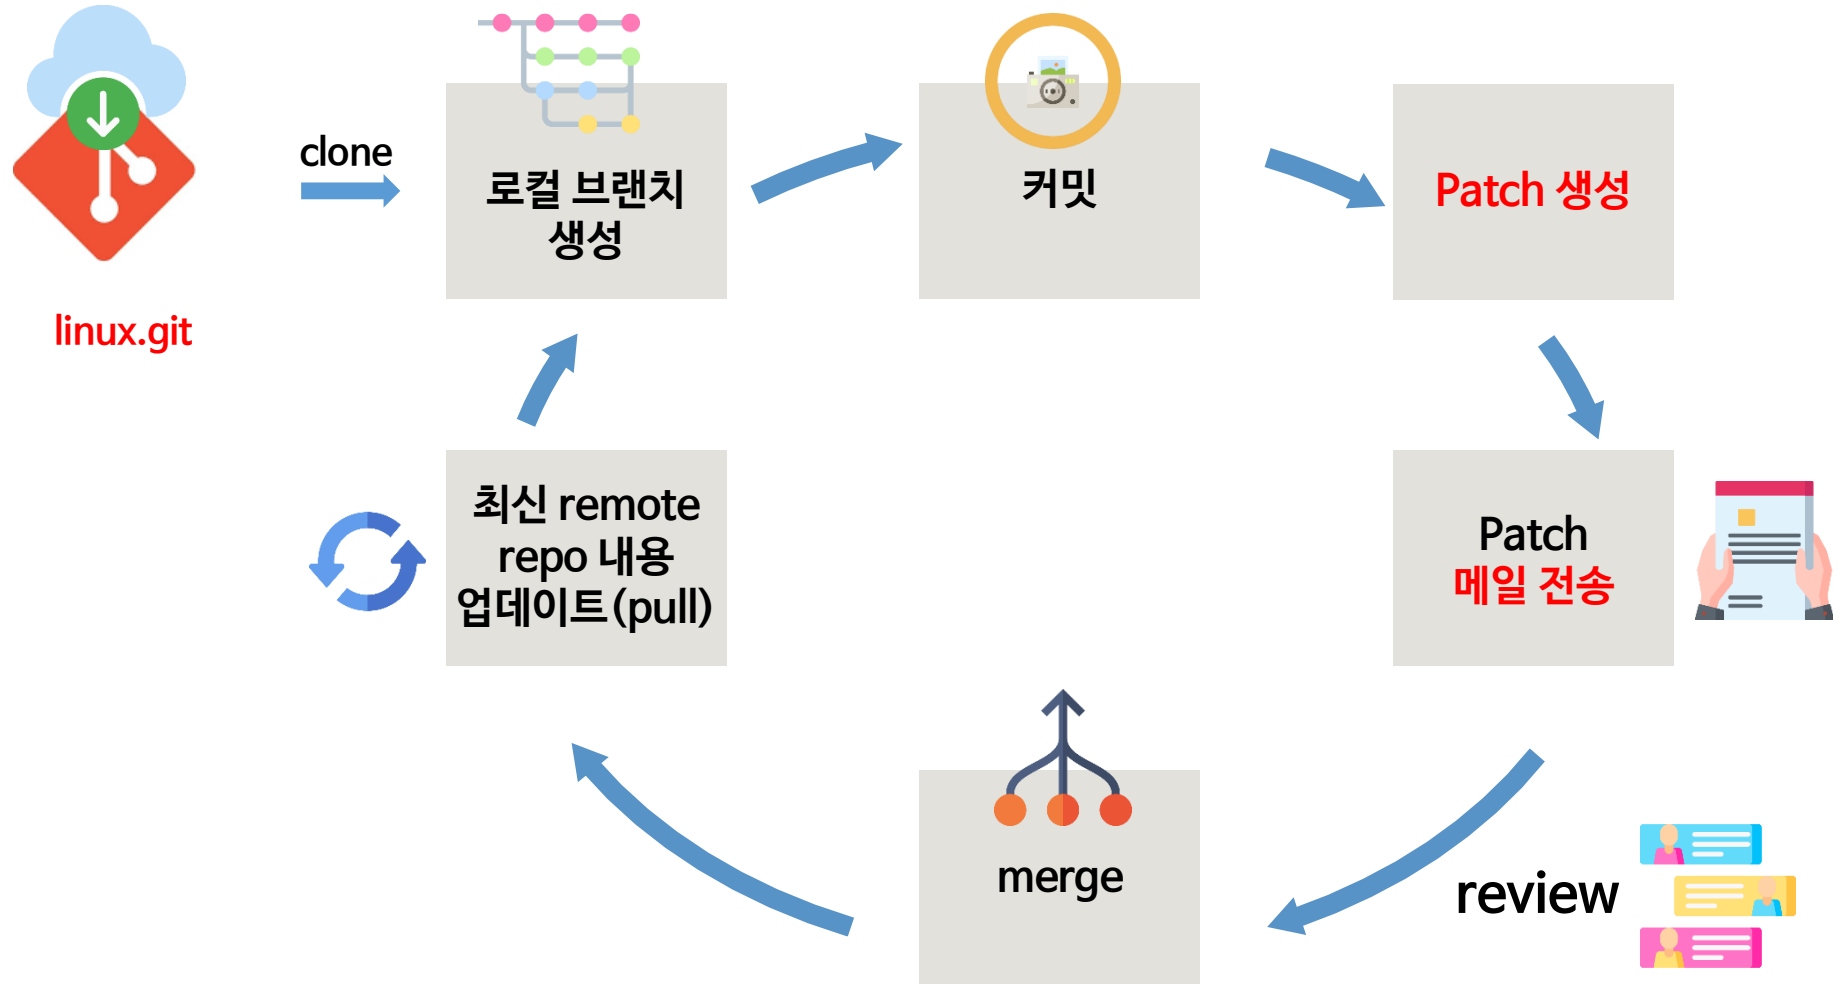
Tag	Download
v5.19-rc5	linux-5.19-rc5.tar.gz
v5.19-rc4	linux-5.19-rc4.tar.gz
v5.19-rc3	linux-5.19-rc3.tar.gz
v5.19-rc2	linux-5.19-rc2.tar.gz
v5.19-rc1	linux-5.19-rc1.tar.gz
v5.18	linux-5.18.tar.gz
v5.18-rc7	linux-5.18-rc7.tar.gz
v5.18-rc6	linux-5.18-rc6.tar.gz
v5.18-rc5	linux-5.18-rc5.tar.gz
v5.18-rc4	linux-5.18-rc4.tar.gz
[...]	

Age	Commit message
8 hours	Merge tag 'loongarch-fixes-5.19-4' of git://git.kernel.org/pub/scm/linux/kern... HEAD master
9 hours	Merge tag 'net-5.19-rc6' of git://git.kernel.org/pub/scm/linux/kernel/git/net...
9 hours	Merge tag 'pinctrl-v5.19-2' of git://git.kernel.org/pub/scm/linux/kernel/git/...
9 hours	signal handling: don't use BUG_ON() for debugging
16 hours	LoongArch: Fix section mismatch warning
16 hours	LoongArch: Fix build errors for tinyconfig
16 hours	LoongArch: Remove obsolete mentions of vcsr
16 hours	LoongArch: Drop these obsolete selects in Kconfig
23 hours	Merge branch 'wireguard-patches-for-5-19-rc6'
23 hours	wireguard: Kconfig: select CRYPTO_CHACHA_S390
[...]	

Linux Kernel 컨트리뷰션 과정?




Linux Kernel 컨트리뷰션 과정?



Linux Kernel 컨트리뷰션?

소스 코드 다운 경로?

메일로 커밋한다고?



패치란 무엇인가?

빌드 환경 구성?

시작하기도 전에 마주한 무수히 많은 장벽들...

이번 세션의 목표

리눅스 커널 컨트리뷰션을 위한 장벽을 허물어 봅시다.

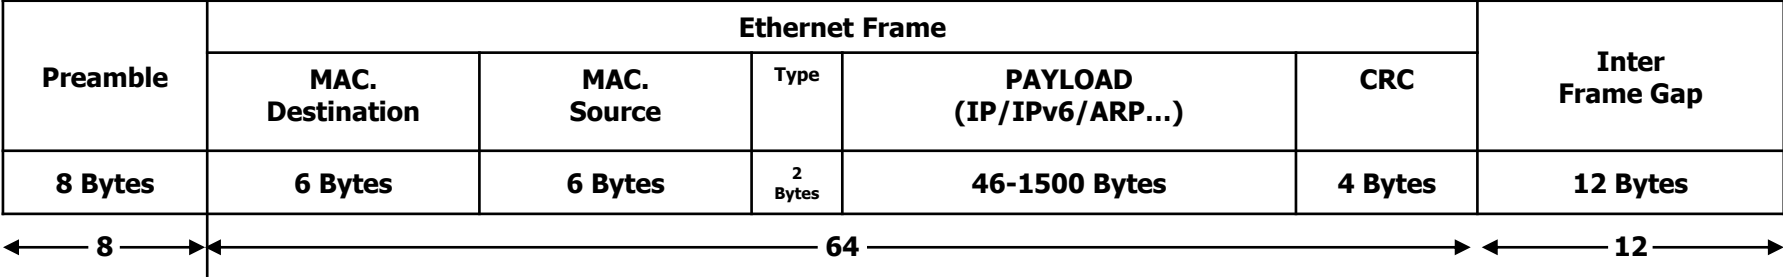

1

Linux Kernel Networking Stack

Packet Path


Linux Kernel Networking Stack

Ethernet Frame




Linux Kernel Networking Stack


**Ethernet
Frame**




**Network
Hardware**



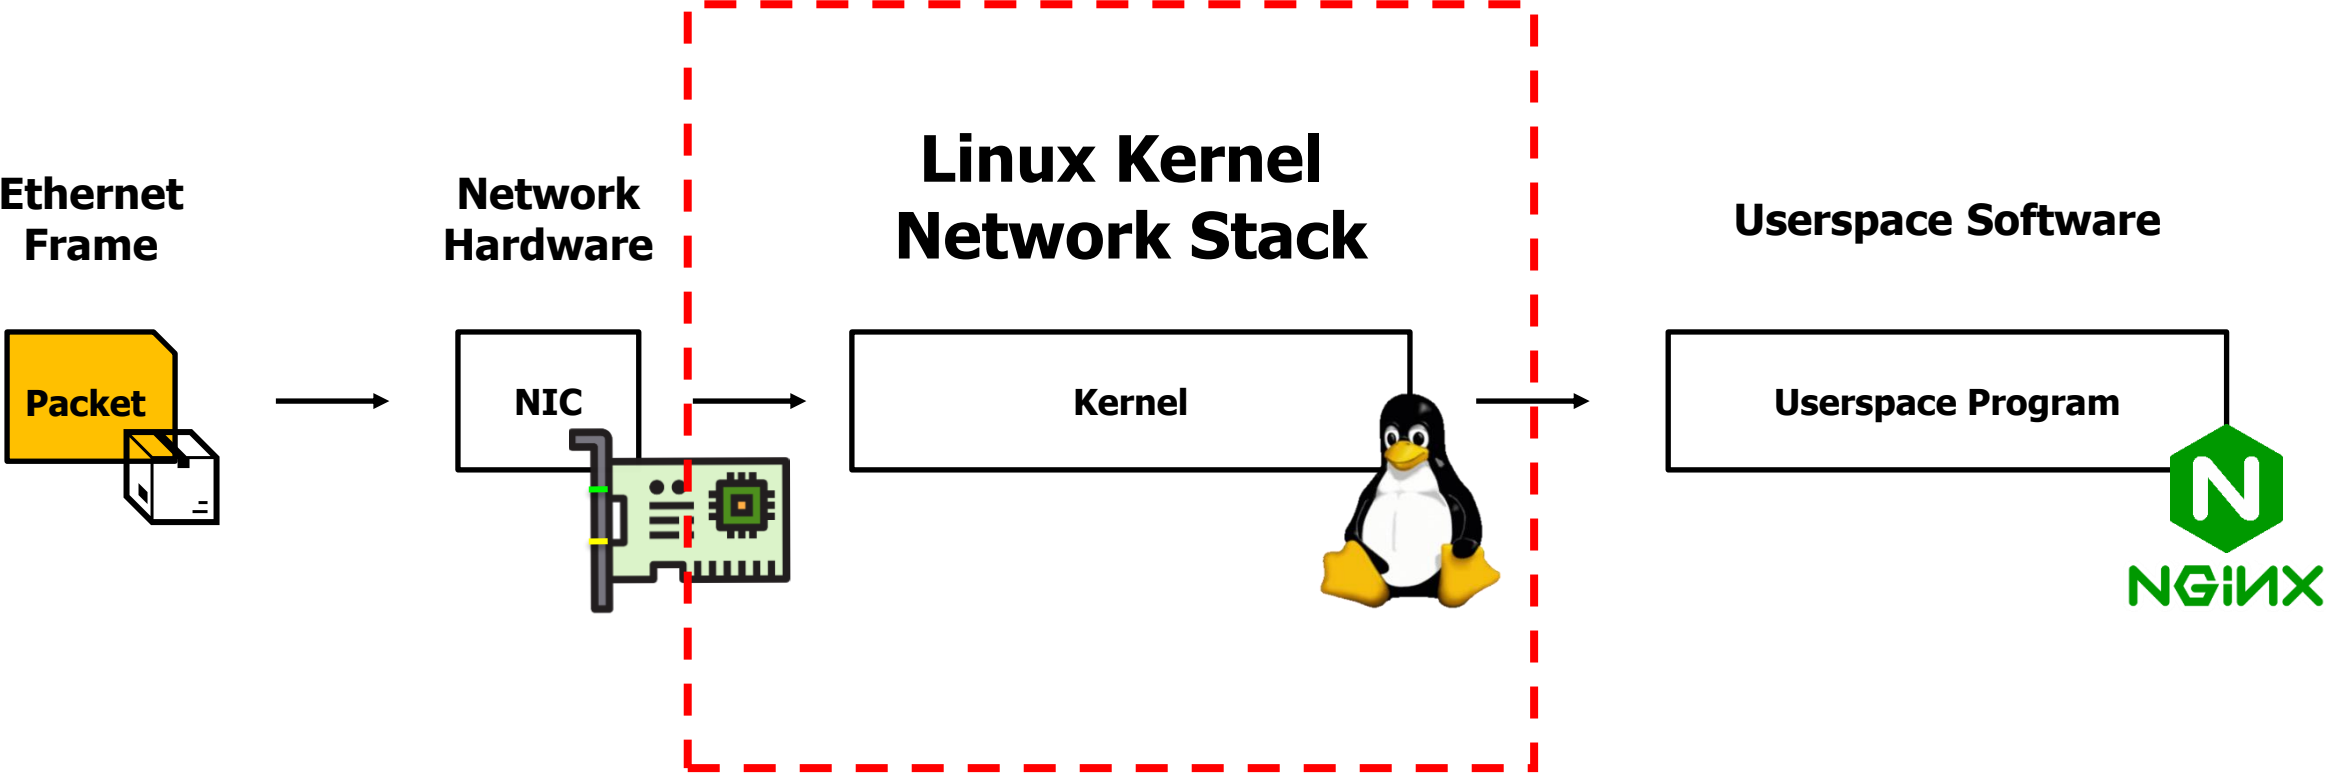
Linux Kernel Networking Stack



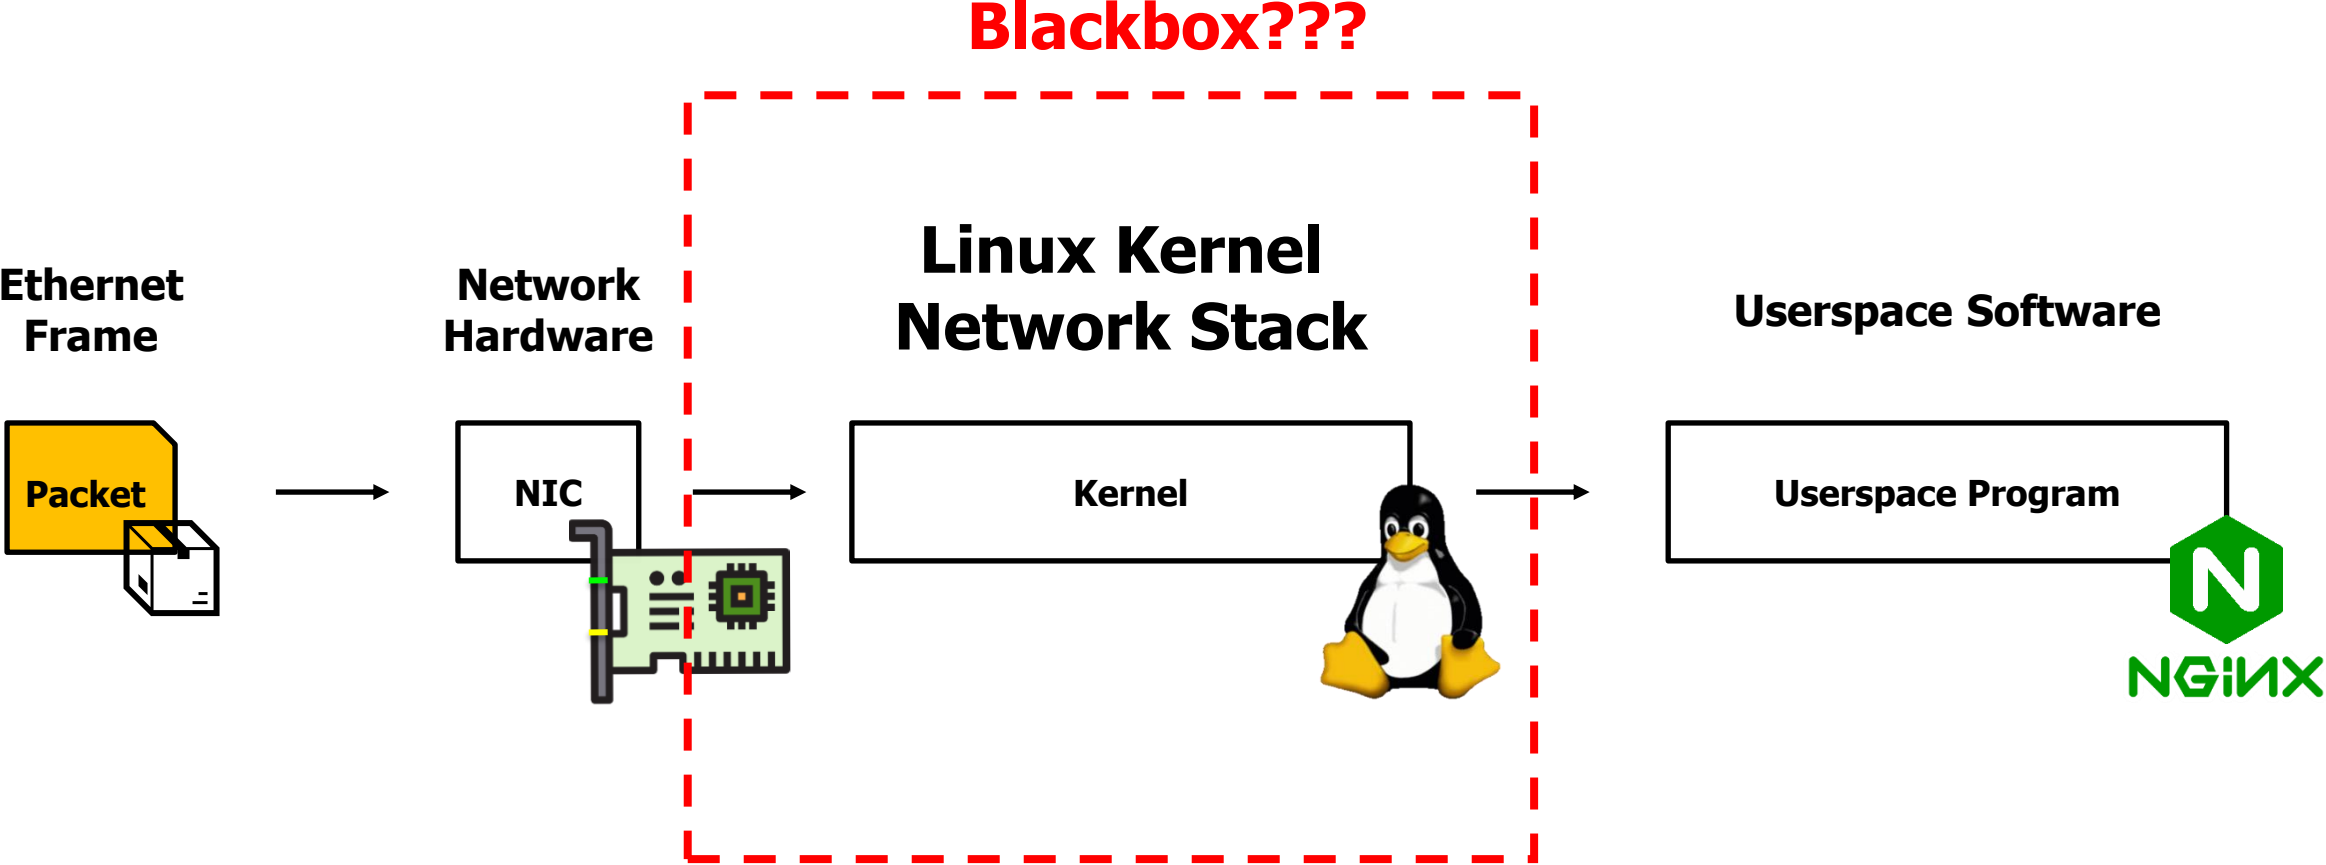
Linux Kernel Networking Stack



Linux Kernel Networking Stack




Linux Kernel Networking Stack




Basic Packet Rx/Tx Path

Packet Path


① Receiving Path




② Forwarding Path



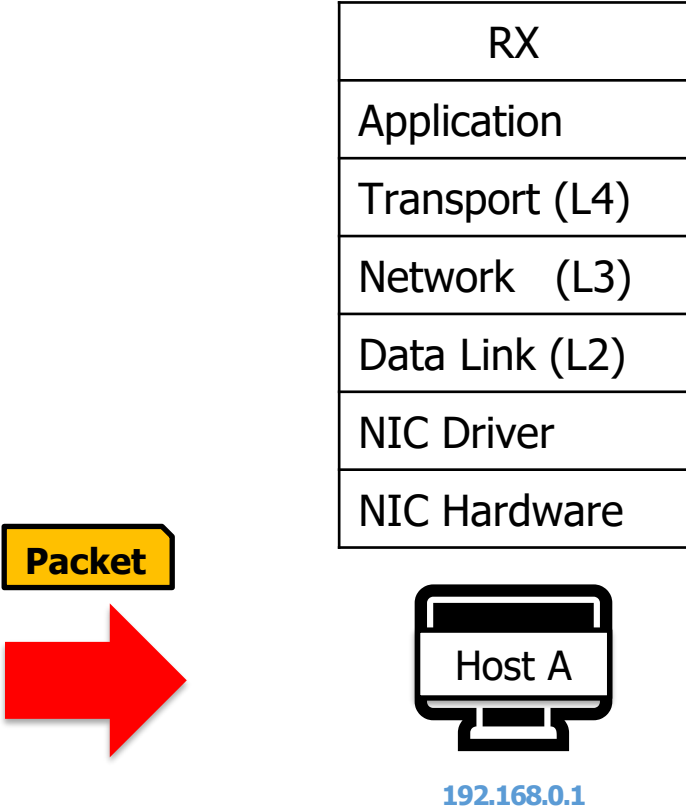
③ Sending Path



Packet Receiving Path

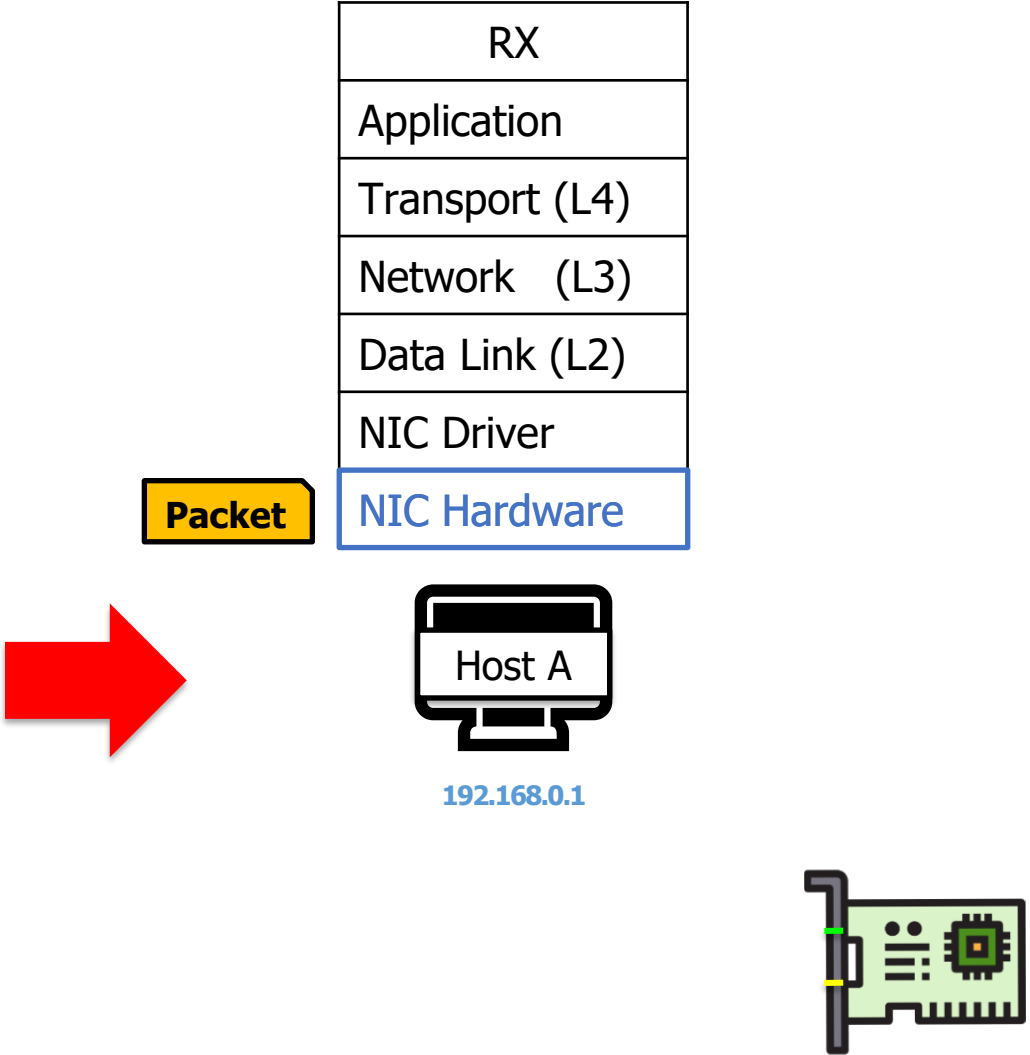


Packet Receiving Path



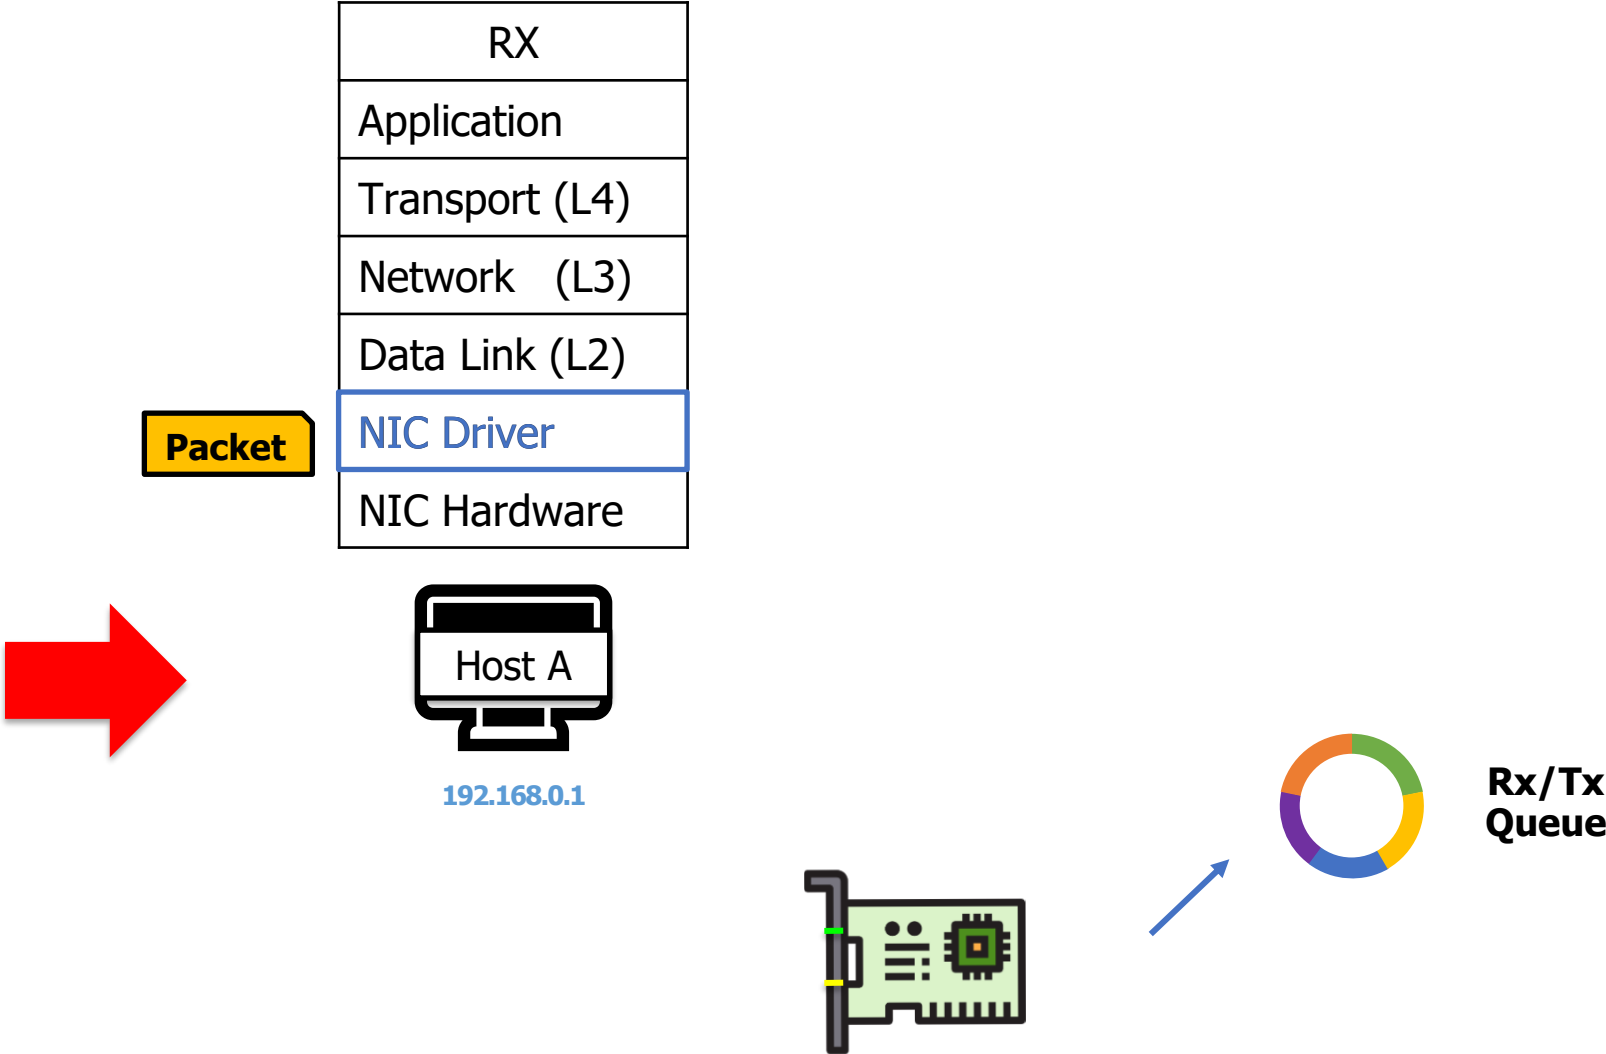
Packet Receiving Path

네트워크 인터페이스 카드 하드웨어에서 패킷을 수신하고



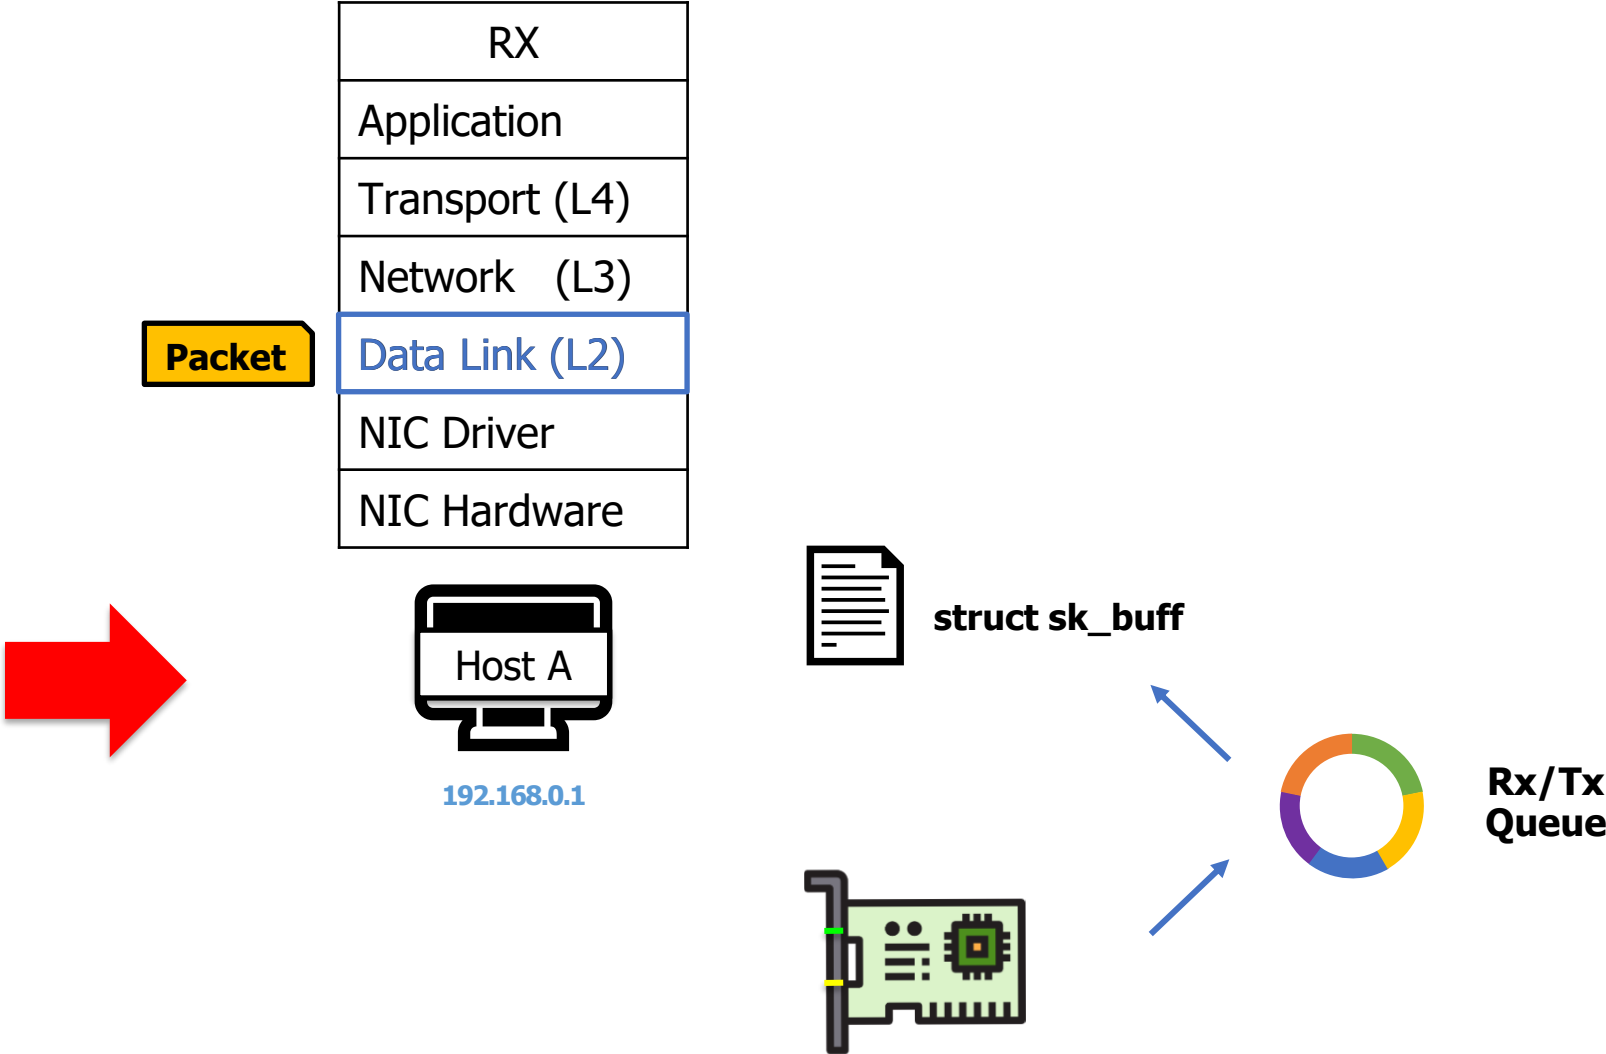
Packet Receiving Path

디바이스 드라이버단에서 NIC의 하드웨어(RX) 큐에서 패킷을 빼 내와서



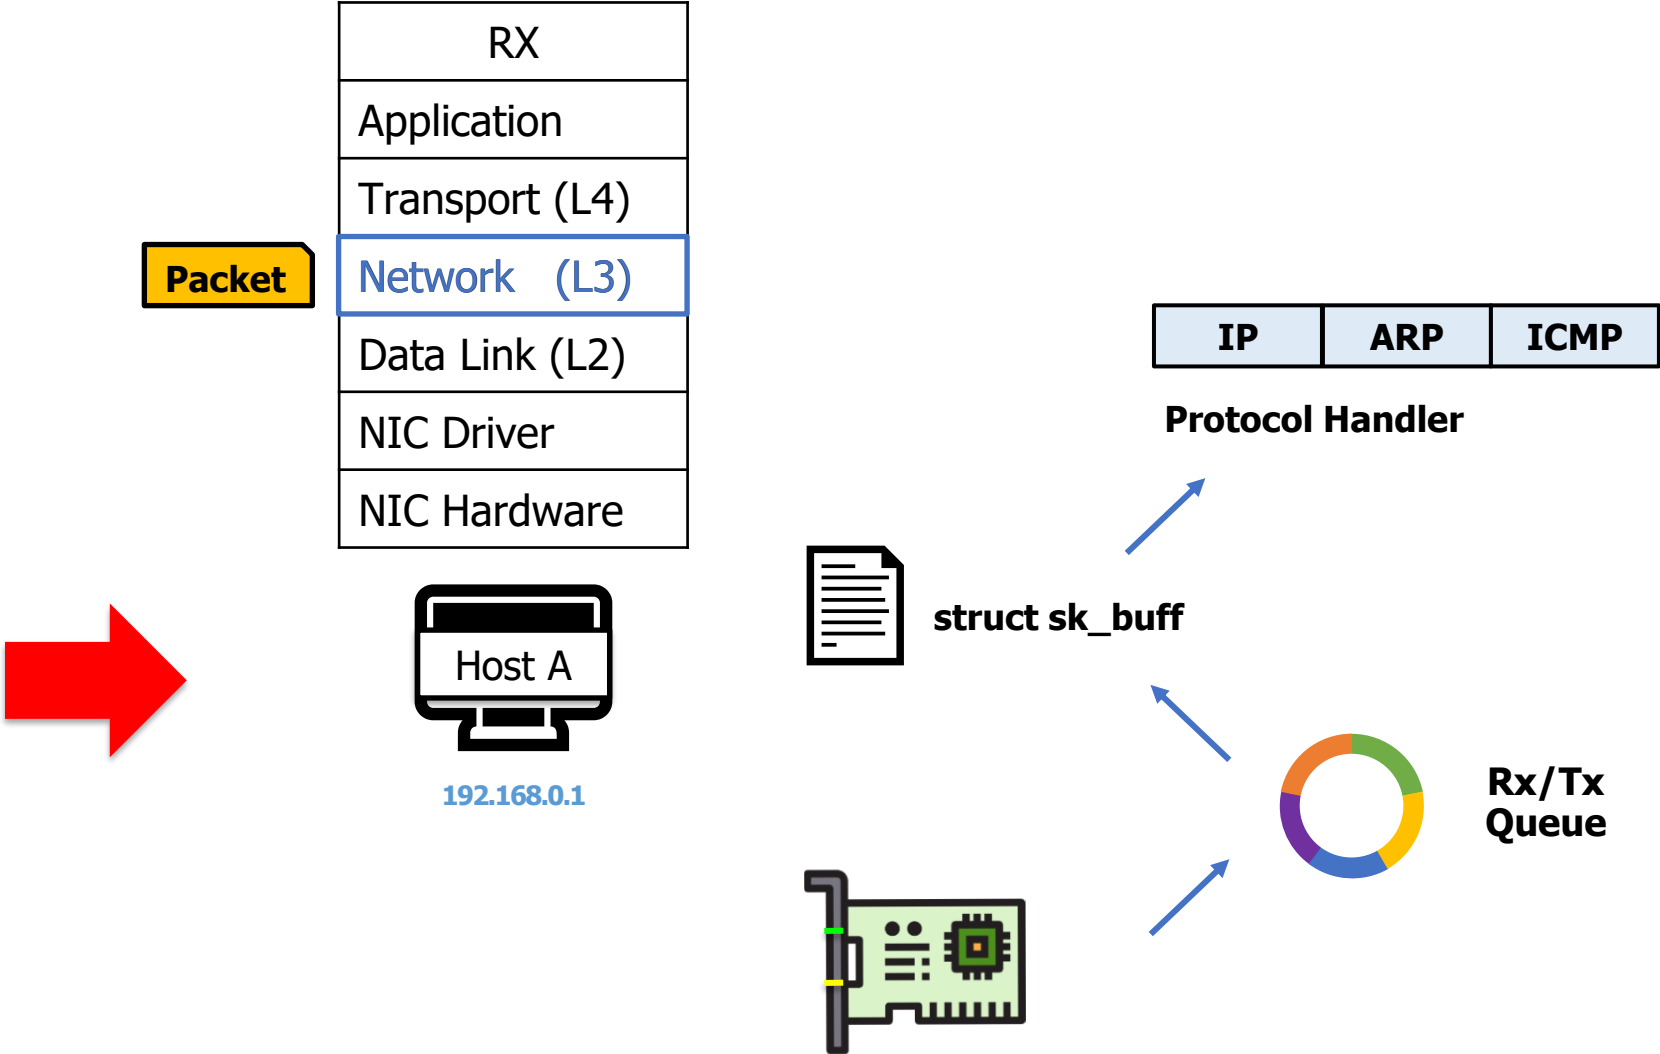
Packet Receiving Path

이를 sk_buff 구조체로 만들고



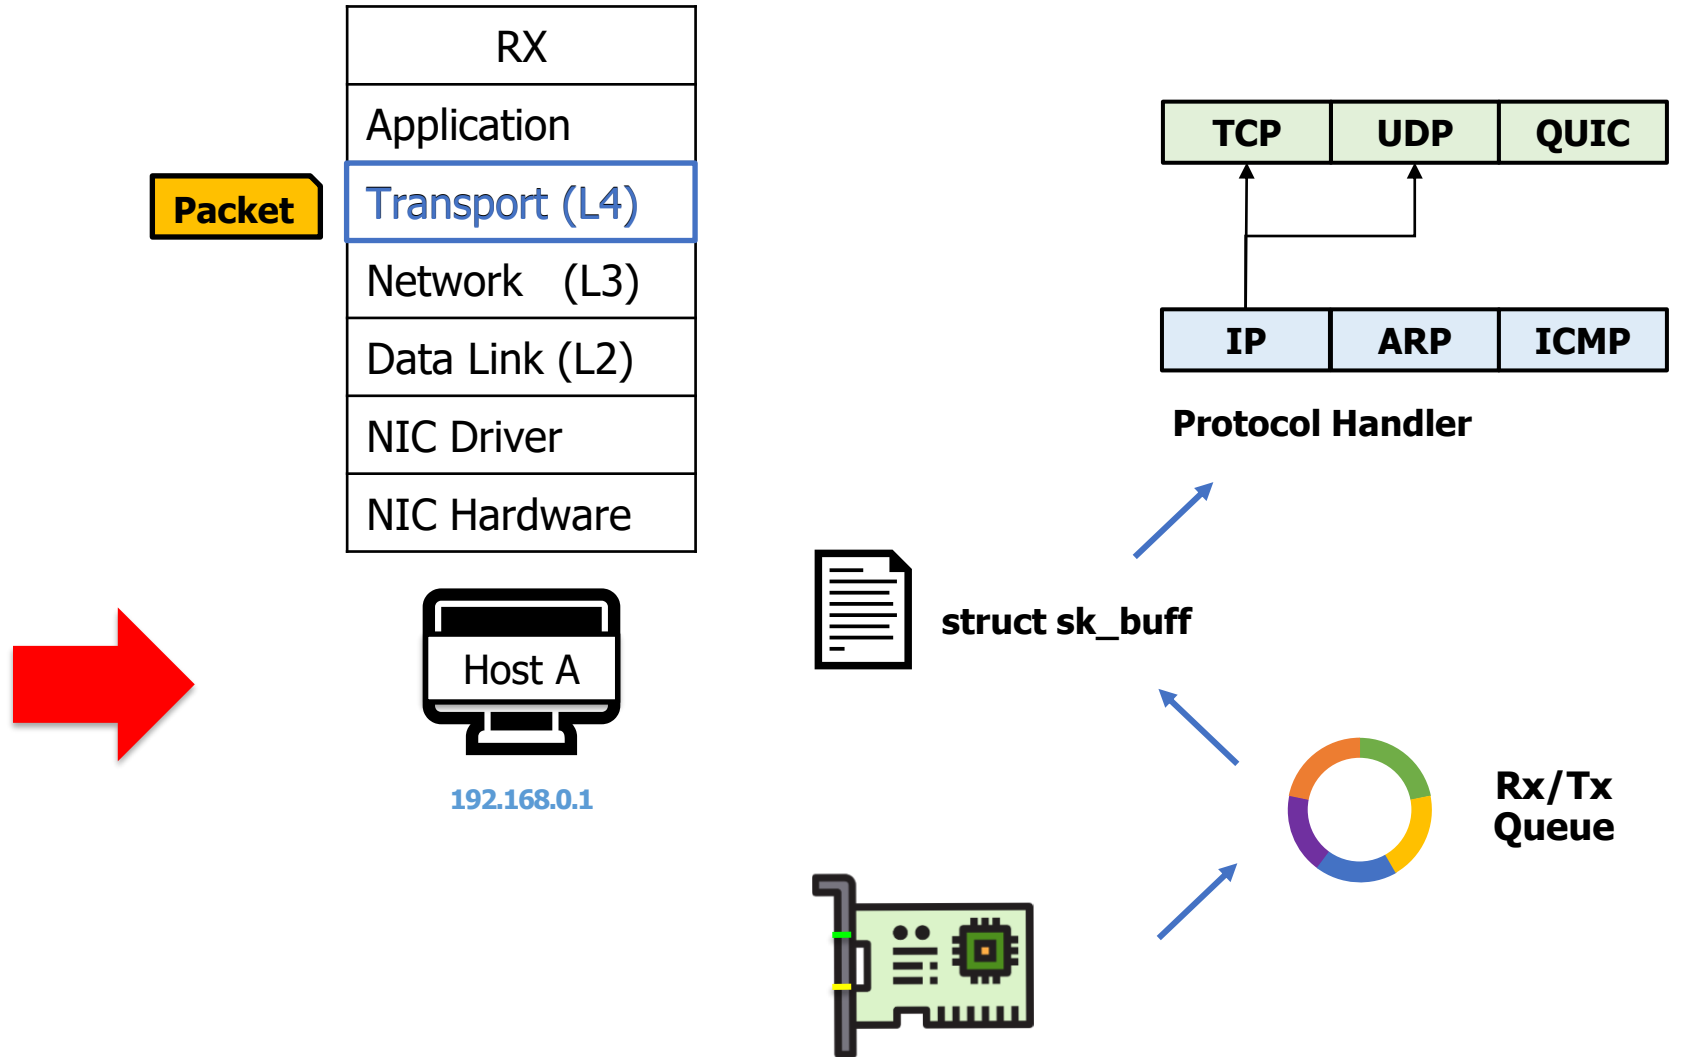
Packet Receiving Path

해당 패킷이 어떤 프로토콜인지 확인하여 해당하는 프로토콜의 핸들러 함수 호출



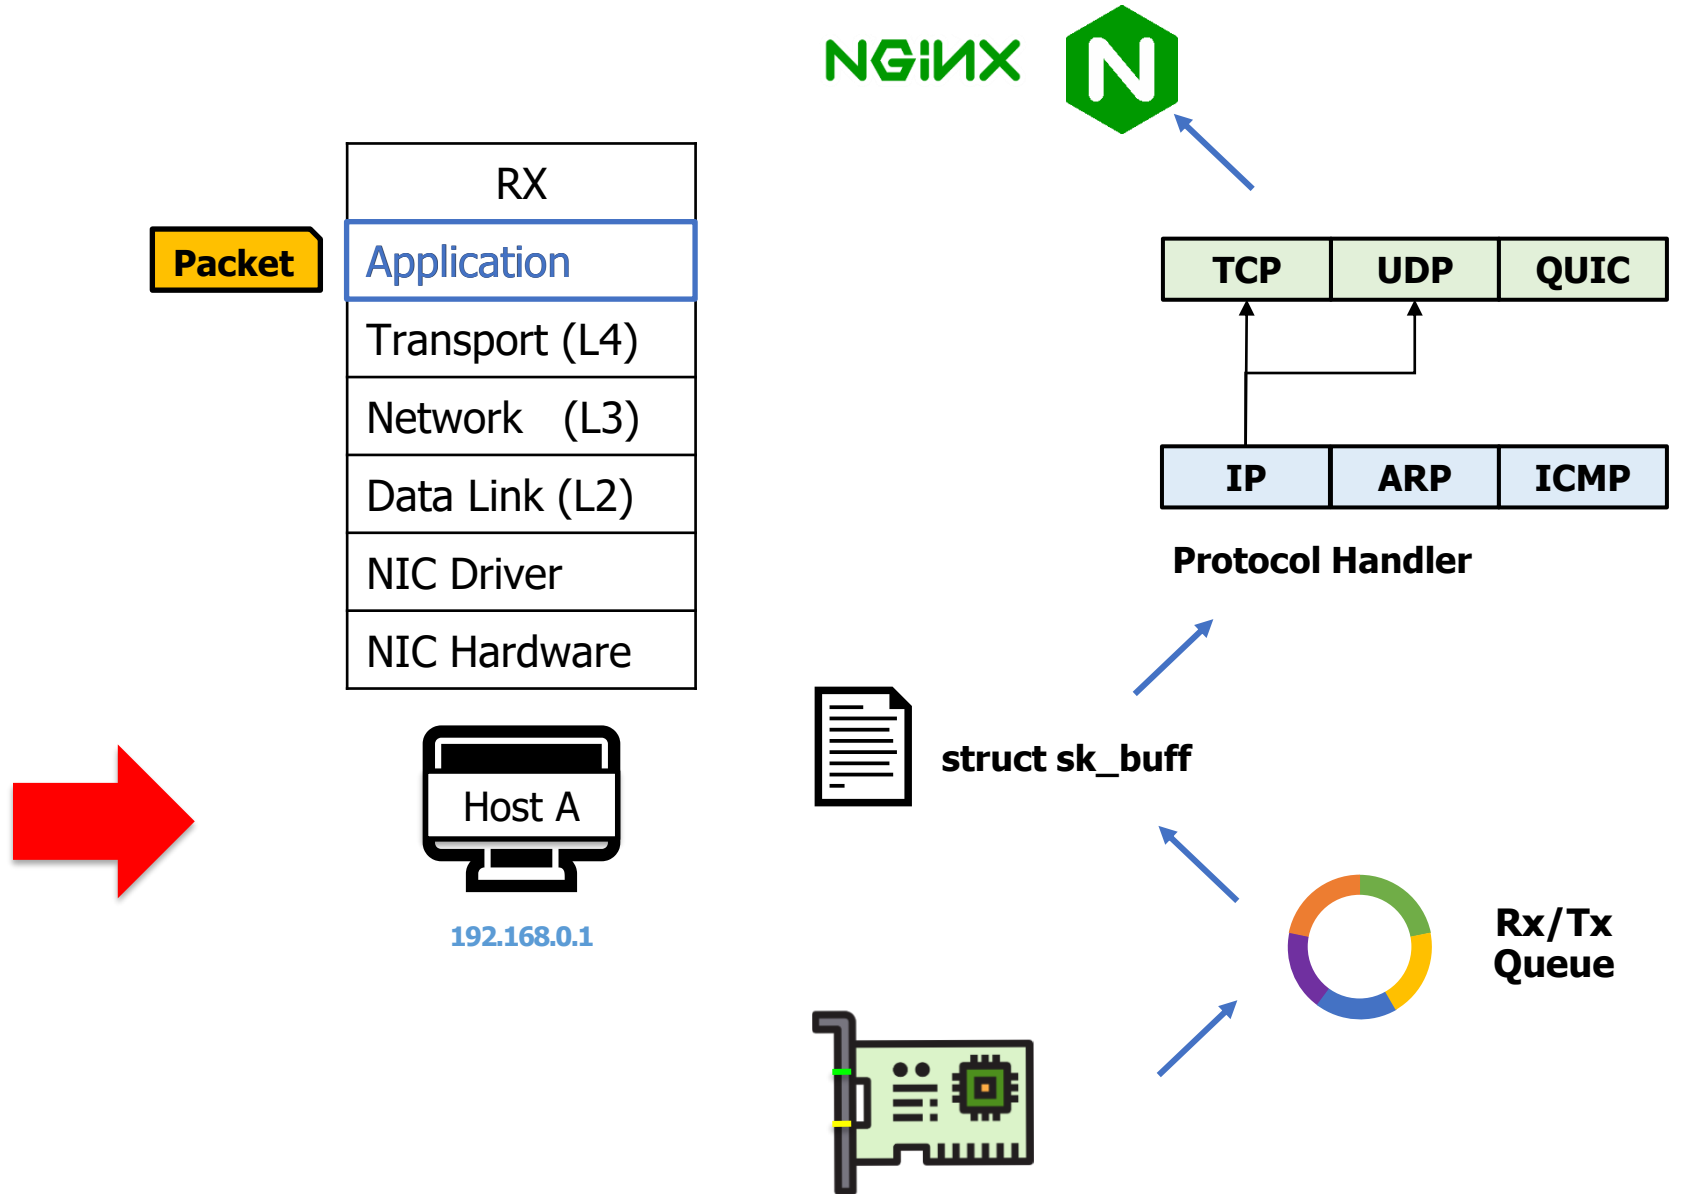
Packet Receiving Path

상위 프로토콜에 따른 적절한 함수를 호출




Packet Receiving Path

처리된 패킷을 사용자 어플리케이션에 넘겨줌



Packet Forwarding Path




Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware



192.168.0.1


Packet



RX/TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware



192.168.0.2/192.168.0.3



RX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.4

Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware



192.168.0.1




RX/TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware



192.168.0.2/192.168.0.3

Packet



RX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.4

Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware



192.168.0.1




Packet

RX/TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware



192.168.0.2/192.168.0.3

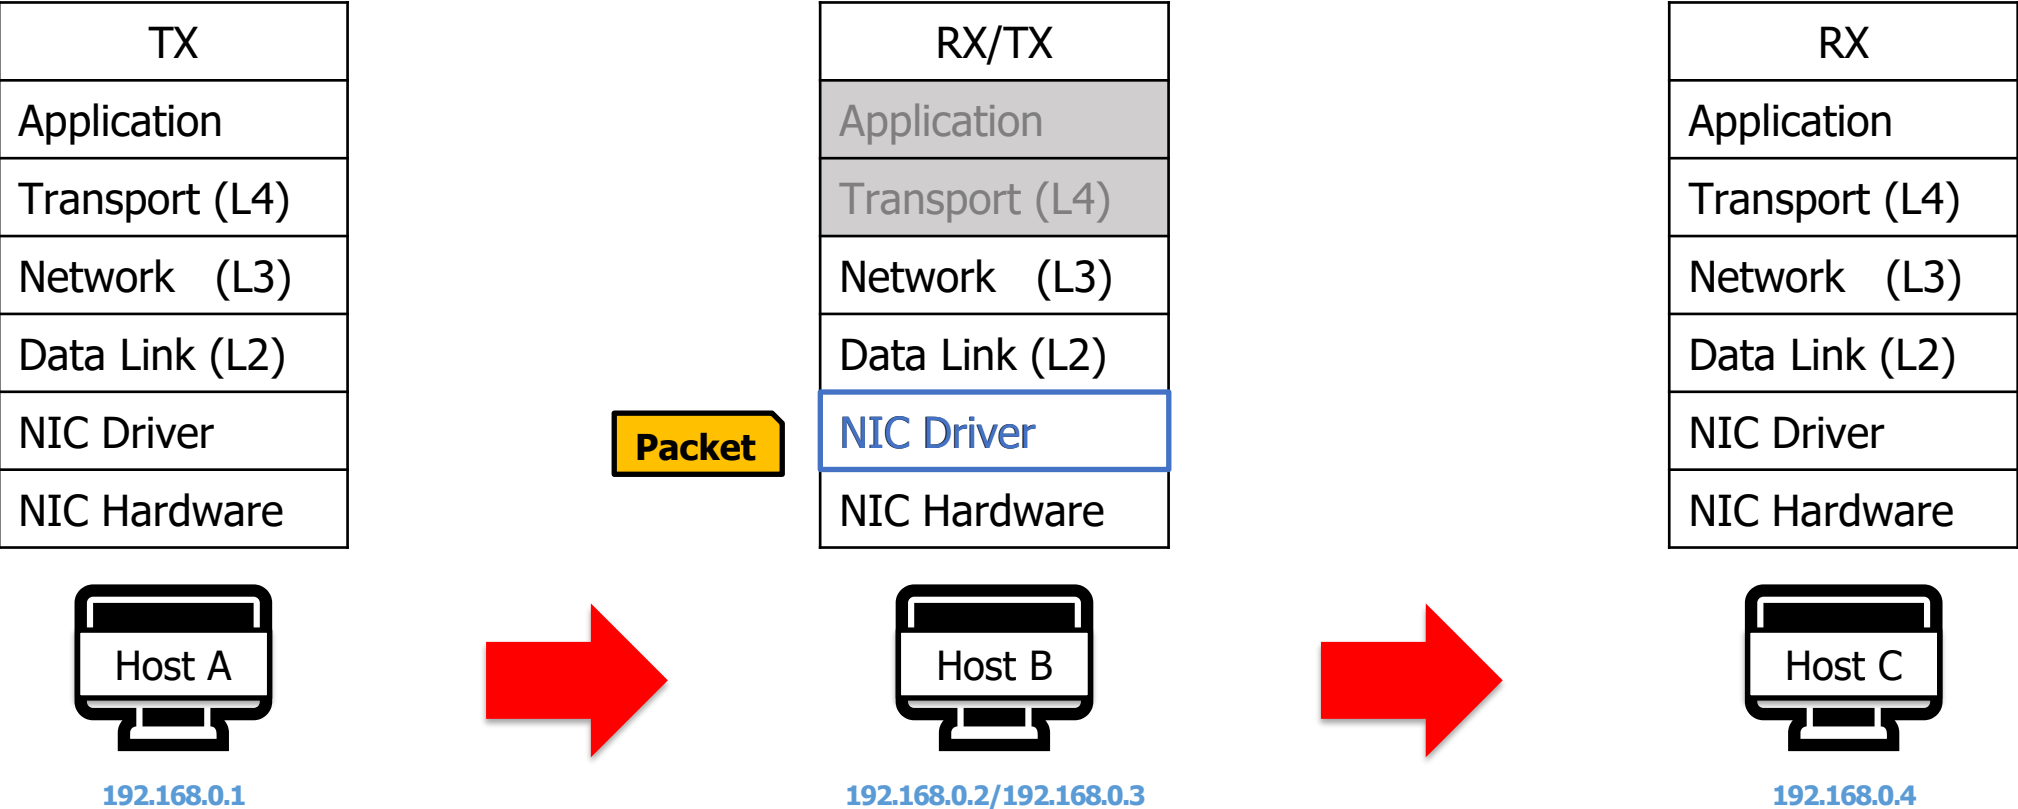


RX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware

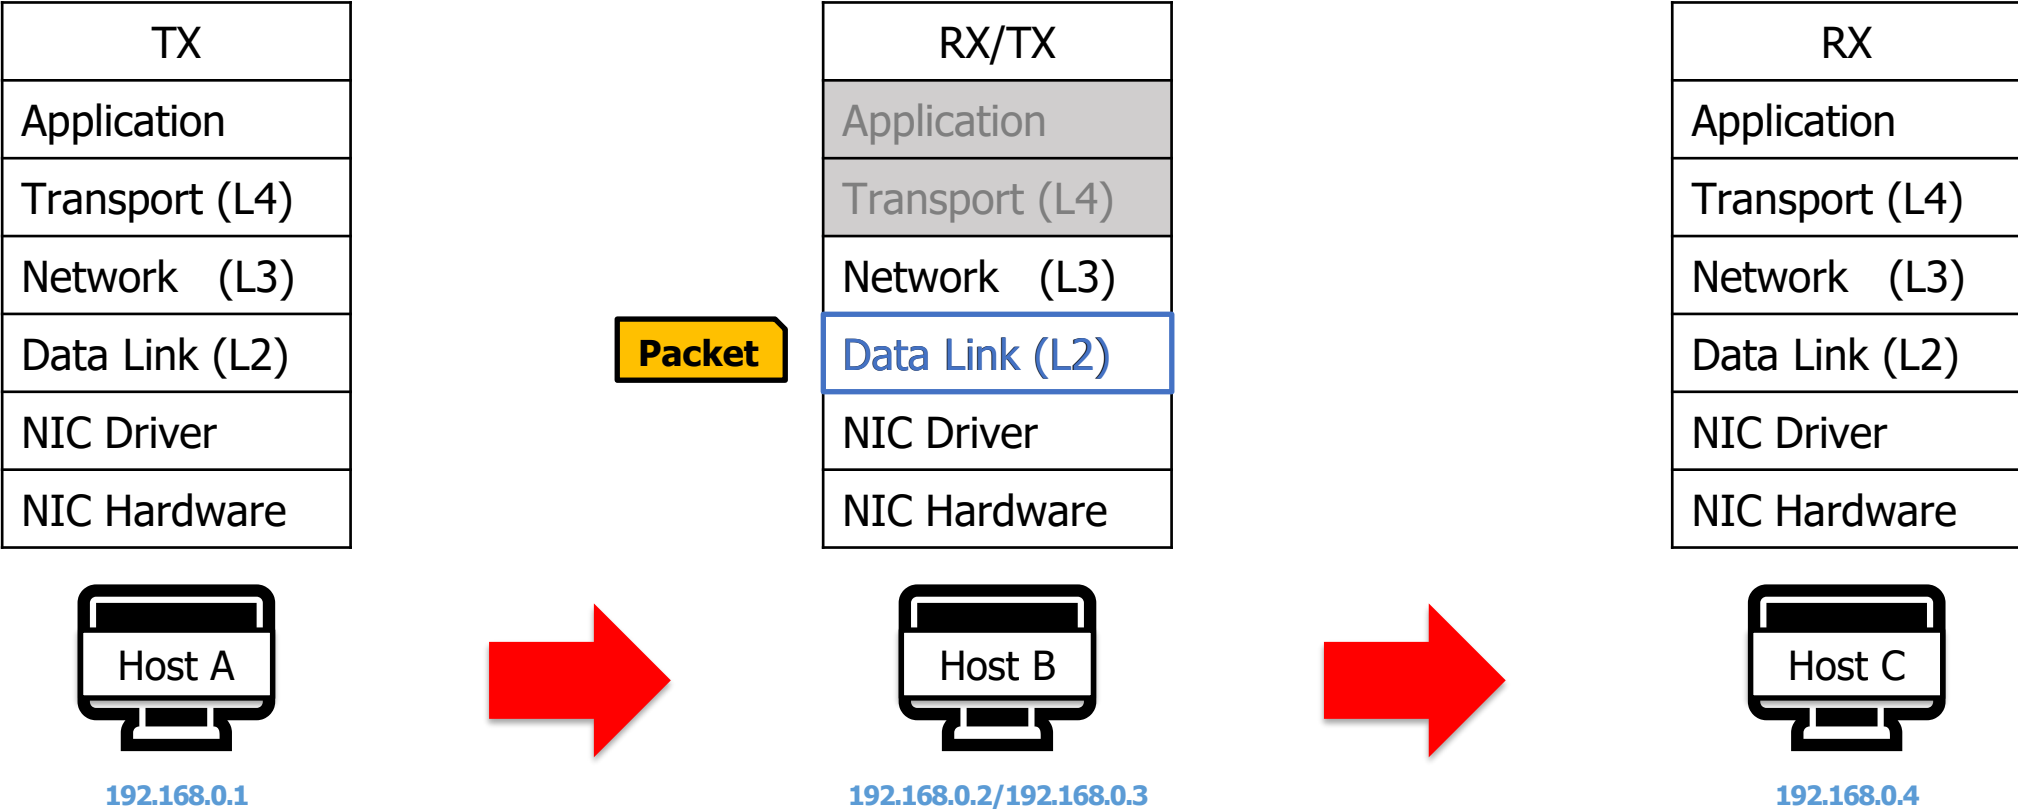


192.168.0.4

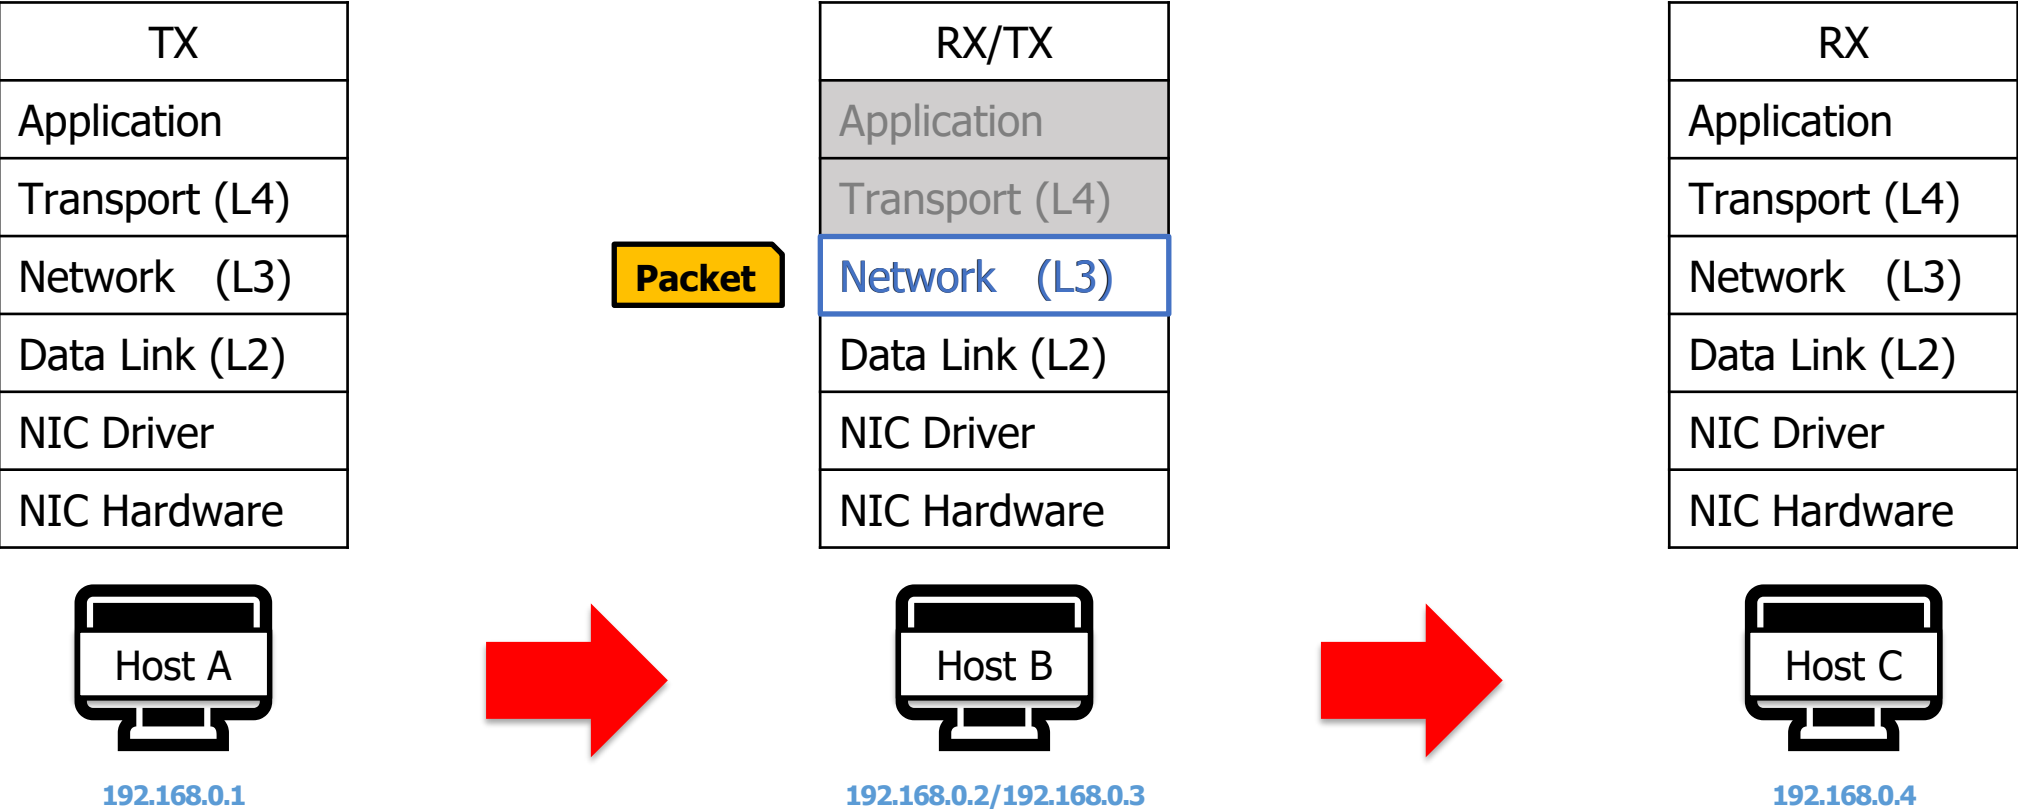
Packet Forwarding Path



Packet Forwarding Path



Packet Forwarding Path



Packet Forwarding Path

```

$ route -n
Kernel IP routing table
Destination      Gateway         Genmask         Flags   Metric  Ref  Use  Iface
0.0.0.0          _gateway       0.0.0.0         UG      101    0    0    eth0
192.168.0.0      0.0.0.0       255.255.255.0  U       101    0    0    eth0
192.168.0.1      192.168.0.2   255.255.255.255 UGH     0      0    0    veth21
192.168.0.4      192.168.0.3   255.255.255.255 UGH     0      0    0    veth23
_gateway         0.0.0.0       255.255.255.255 UH      101    0    0    eth0
    
```

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware


Packet

Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware


Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware



192.168.0.1



192.168.0.2/192.168.0.3



192.168.0.4

Packet Forwarding Path

```

$ route -n
Kernel IP routing table
Destination      Gateway         Genmask         Flags   Metric  Ref  Use  Iface
0.0.0.0          _gateway       0.0.0.0         UG      101    0    0    eth0
192.168.0.0      0.0.0.0        255.255.255.0  U       101    0    0    eth0
192.168.0.1      192.168.0.2   255.255.255.255 UGH     0     0    0    veth21
192.168.0.4      192.168.0.3   255.255.255.255 UGH     0     0    0    veth23
_gateway         0.0.0.0        255.255.255.255 UH      101    0    0    eth0
    
```

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware


Packet

Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware


Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware



192.168.0.1

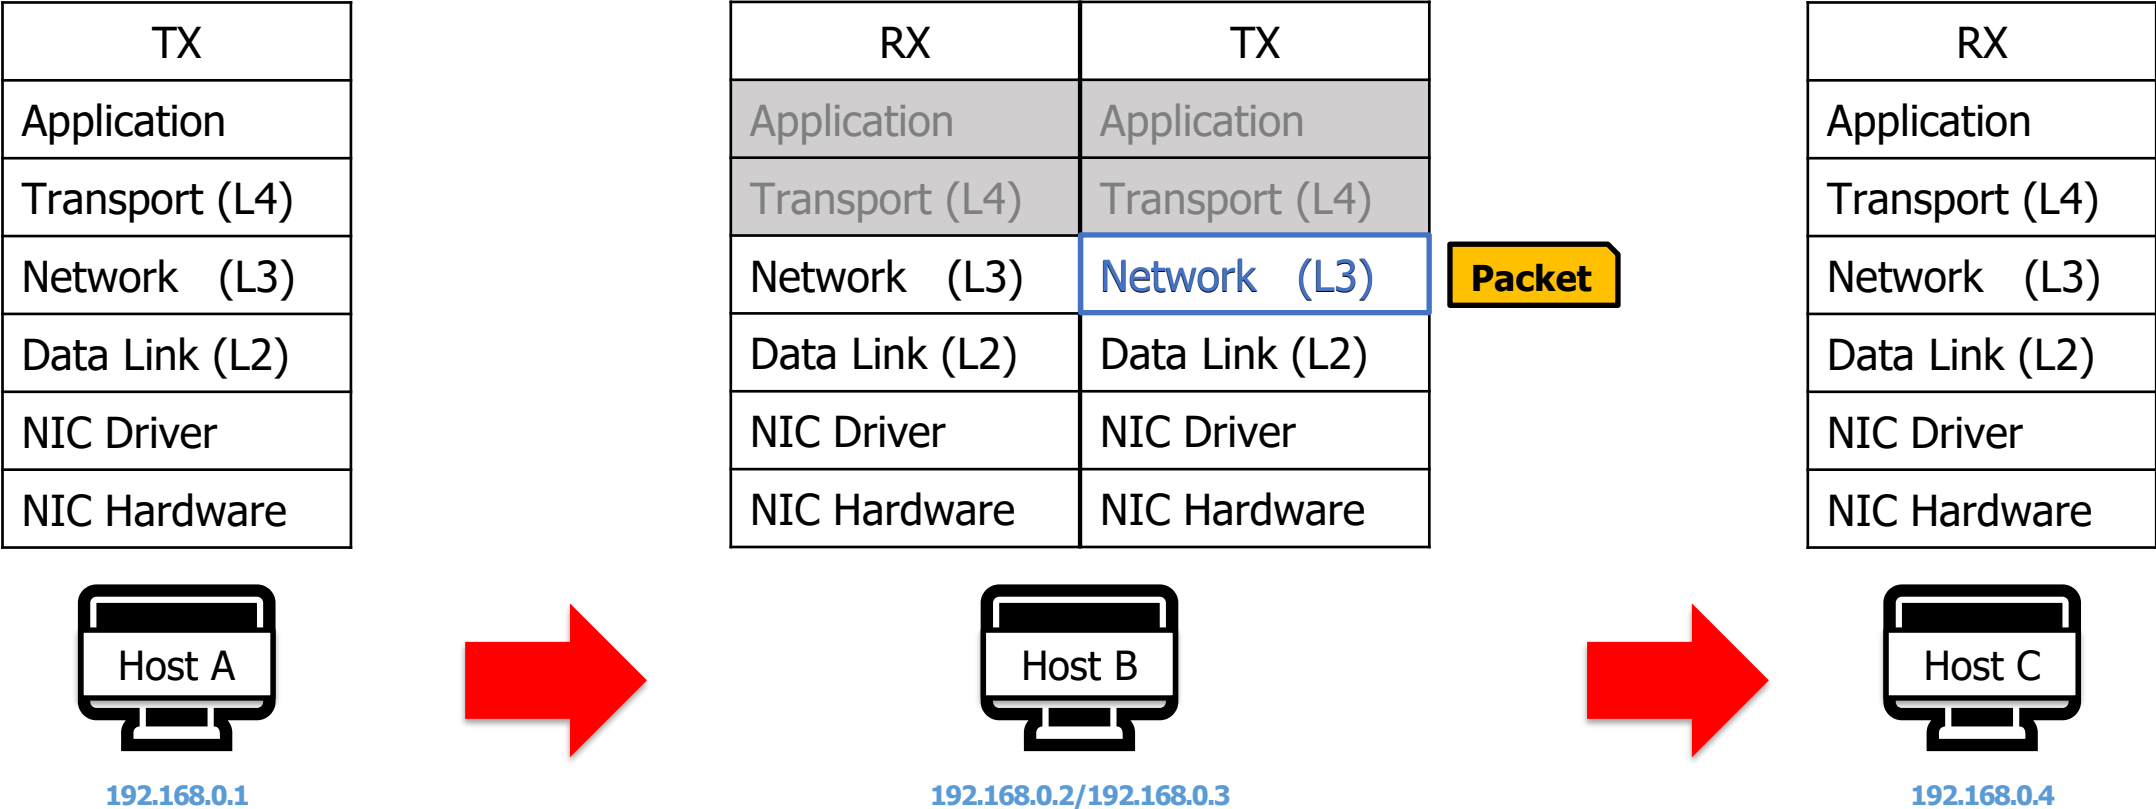


192.168.0.2/192.168.0.3

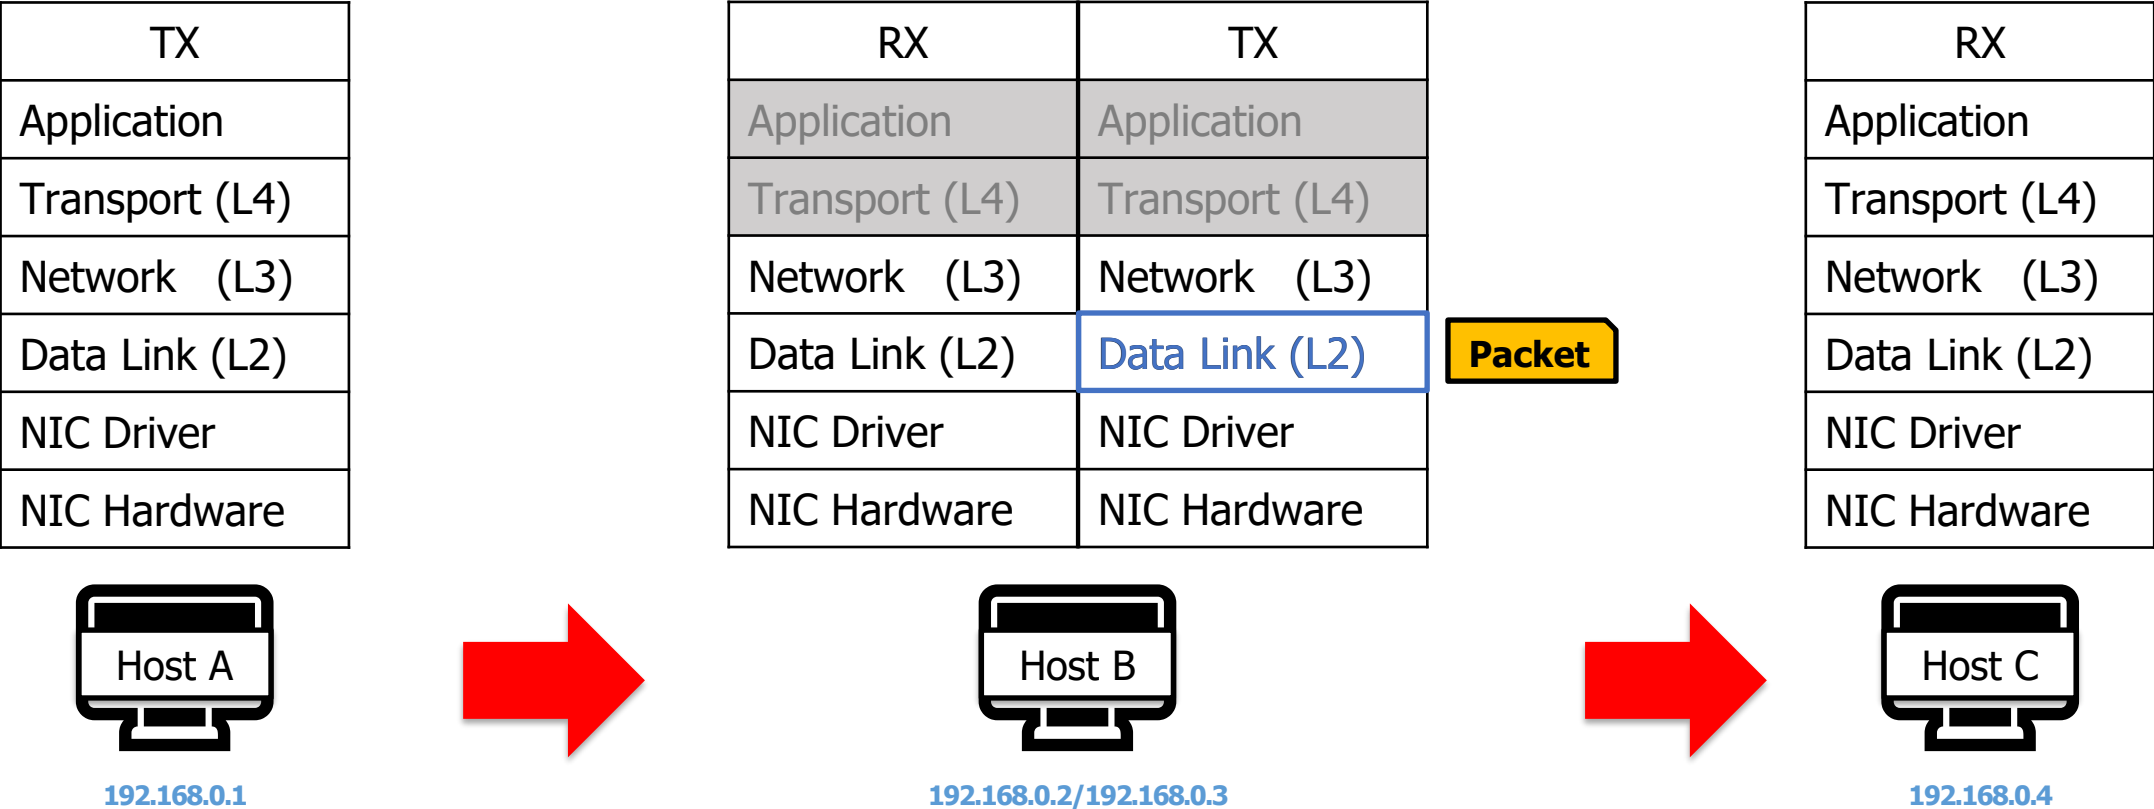


192.168.0.4

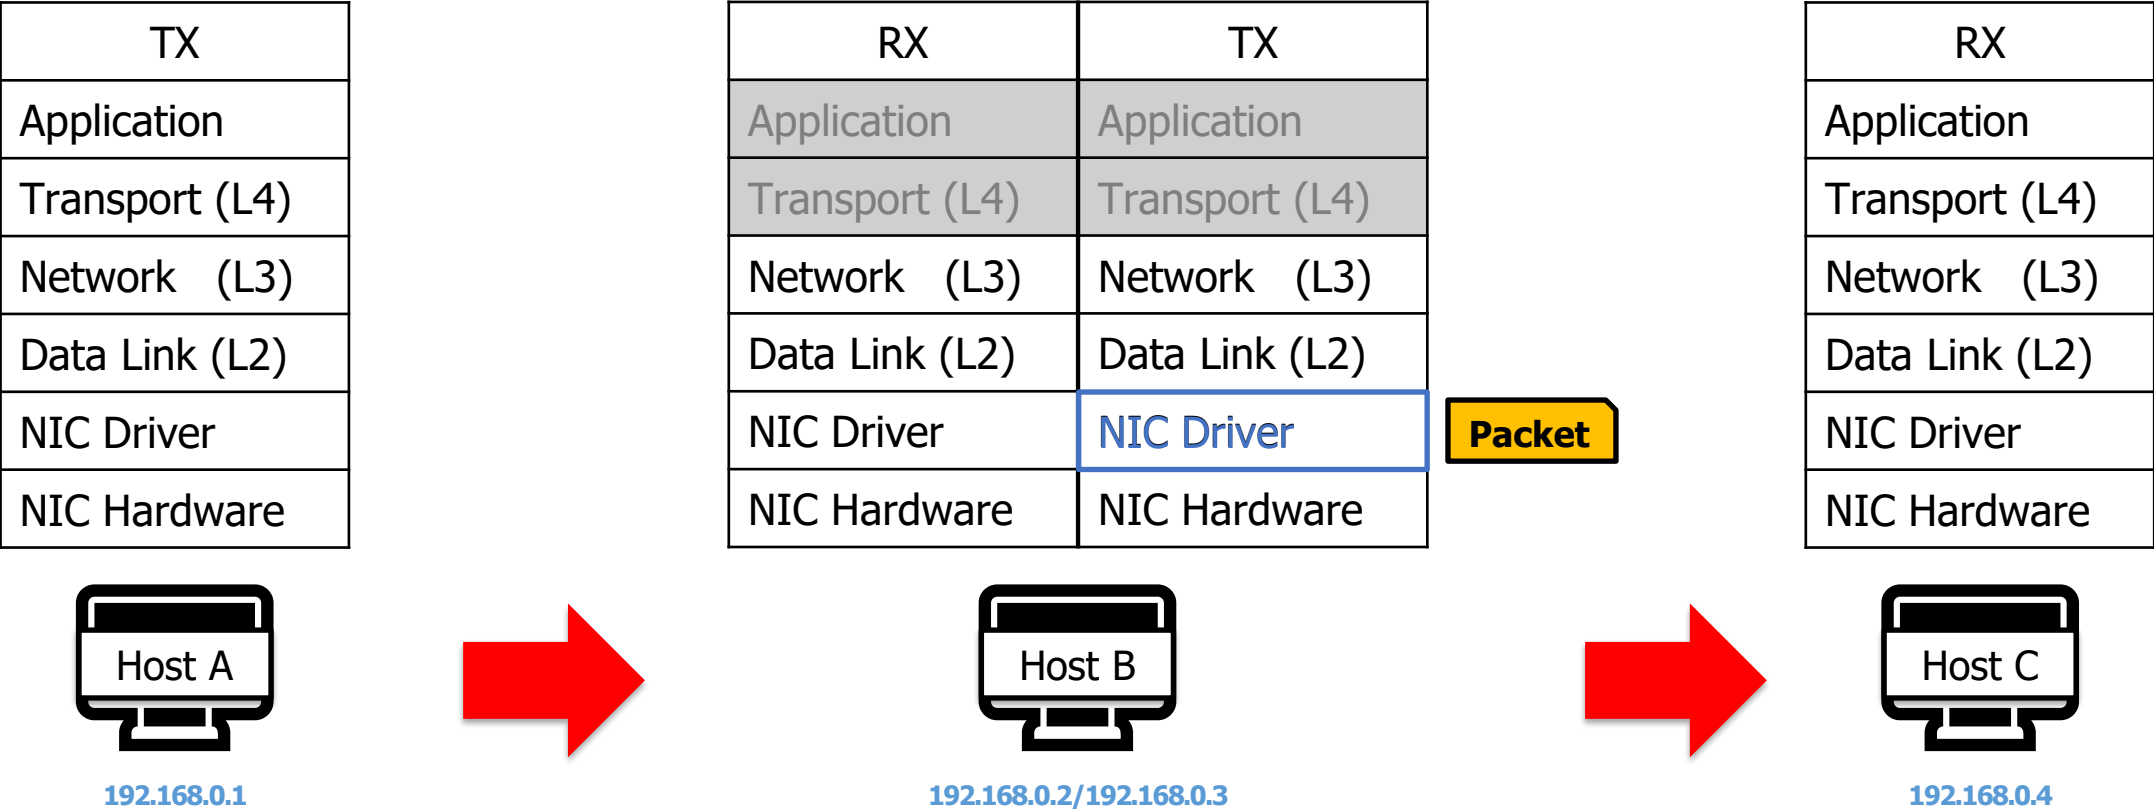
Packet Forwarding Path



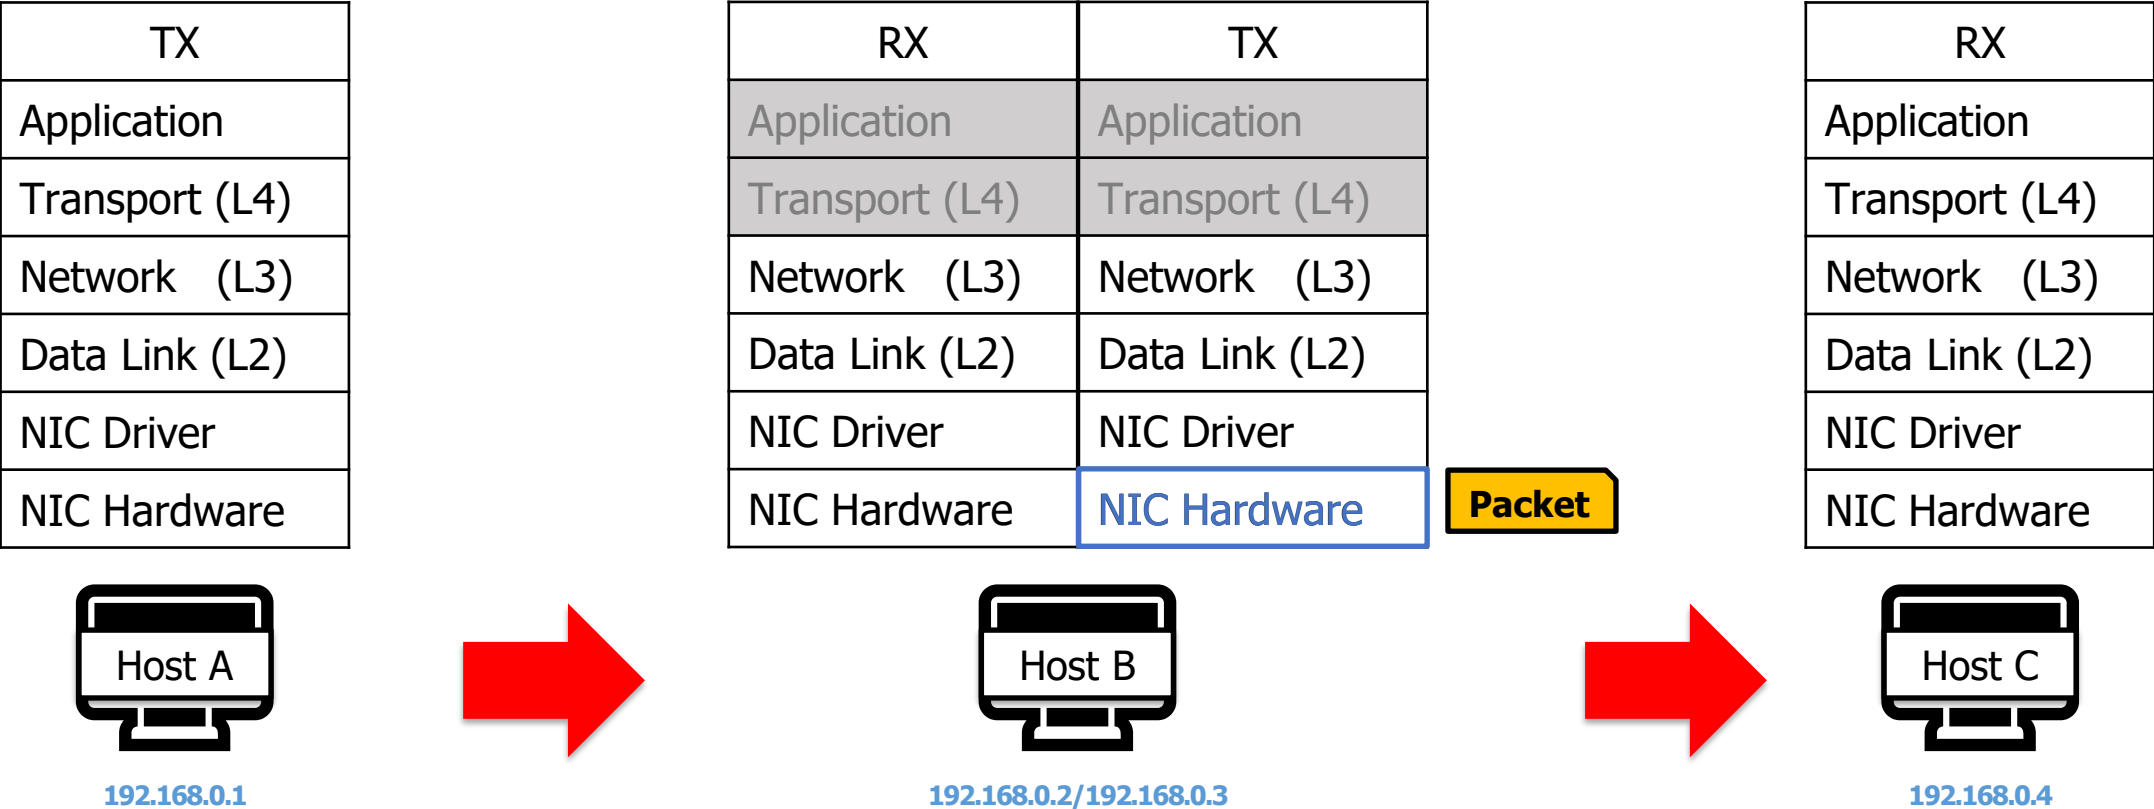
Packet Forwarding Path



Packet Forwarding Path



Packet Forwarding Path




Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.1



RX	TX
Application	Application
Transport (L4)	Transport (L4)
Network (L3)	Network (L3)
Data Link (L2)	Data Link (L2)
NIC Driver	NIC Driver
NIC Hardware	NIC Hardware



192.168.0.2/192.168.0.3



RX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.4

Packet Forwarding Path

TX	
Application	
Transport (L4)	
Network (L3)	
Data Link (L2)	
NIC Driver	
NIC Hardware	




192.168.0.1



RX		TX	
Application		Application	
Transport (L4)		Transport (L4)	
Network (L3)		Network (L3)	
Data Link (L2)		Data Link (L2)	
NIC Driver		NIC Driver	
NIC Hardware		NIC Hardware	



192.168.0.2/192.168.0.3



RX	
Application	
Transport (L4)	
Network (L3)	
Data Link (L2)	
NIC Driver	
Packet	NIC Hardware




192.168.0.4

Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.1



RX	TX
Application	Application
Transport (L4)	Transport (L4)
Network (L3)	Network (L3)
Data Link (L2)	Data Link (L2)
NIC Driver	NIC Driver
NIC Hardware	NIC Hardware



192.168.0.2/192.168.0.3



Packet

RX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.4

Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.1



RX	TX
Application	Application
Transport (L4)	Transport (L4)
Network (L3)	Network (L3)
Data Link (L2)	Data Link (L2)
NIC Driver	NIC Driver
NIC Hardware	NIC Hardware



192.168.0.2/192.168.0.3



Packet

RX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.4

Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.1



RX	TX
Application	Application
Transport (L4)	Transport (L4)
Network (L3)	Network (L3)
Data Link (L2)	Data Link (L2)
NIC Driver	NIC Driver
NIC Hardware	NIC Hardware



192.168.0.2/192.168.0.3



Packet

RX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.4

Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.1



RX	TX
Application	Application
Transport (L4)	Transport (L4)
Network (L3)	Network (L3)
Data Link (L2)	Data Link (L2)
NIC Driver	NIC Driver
NIC Hardware	NIC Hardware



192.168.0.2/192.168.0.3



Packet

RX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.4

Packet Forwarding Path

TX
Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware




192.168.0.1



RX	TX
Application	Application
Transport (L4)	Transport (L4)
Network (L3)	Network (L3)
Data Link (L2)	Data Link (L2)
NIC Driver	NIC Driver
NIC Hardware	NIC Hardware



192.168.0.2/192.168.0.3




RX
Packet Application
Transport (L4)
Network (L3)
Data Link (L2)
NIC Driver
NIC Hardware

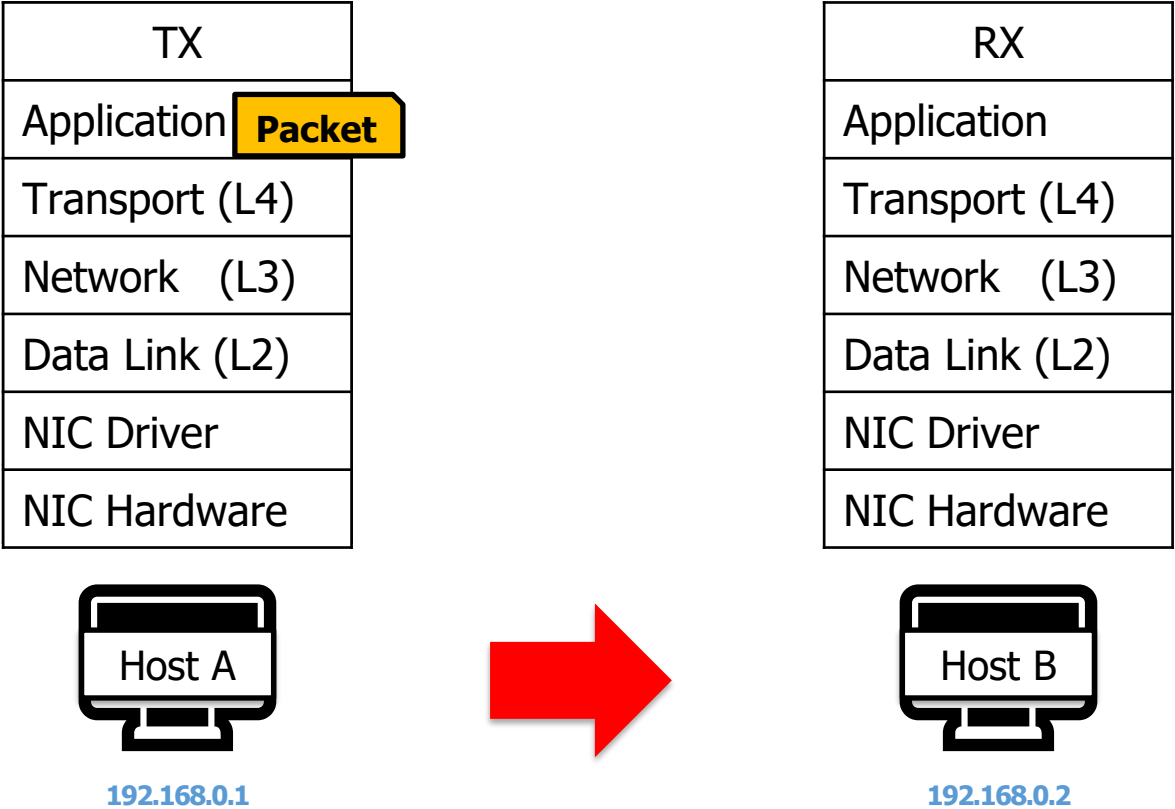


192.168.0.4

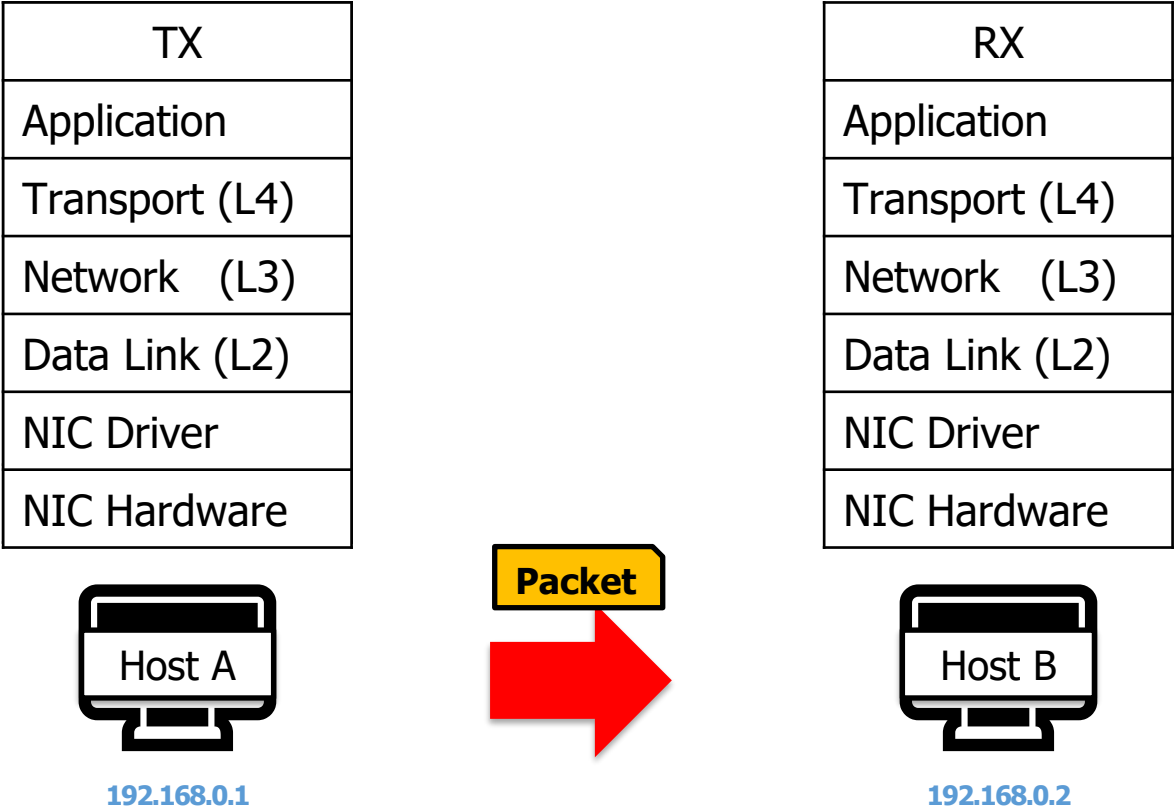
Packet Sending Path



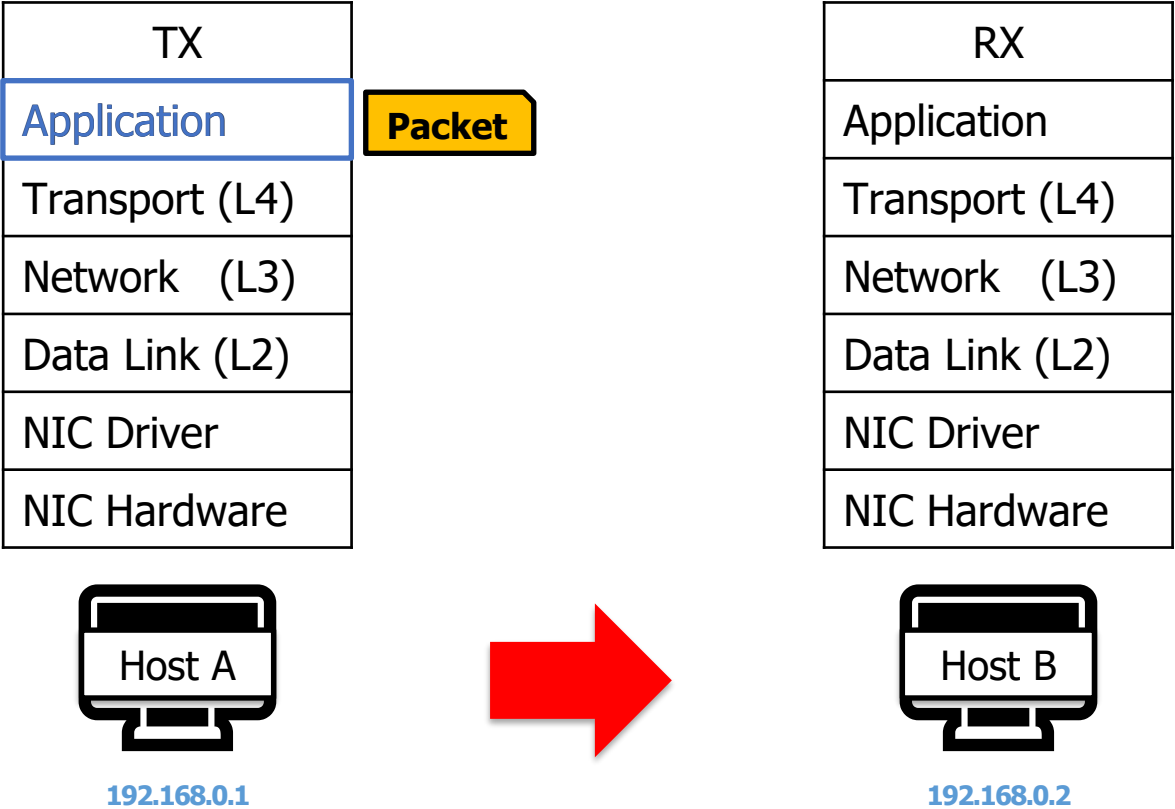
Packet Sending Path



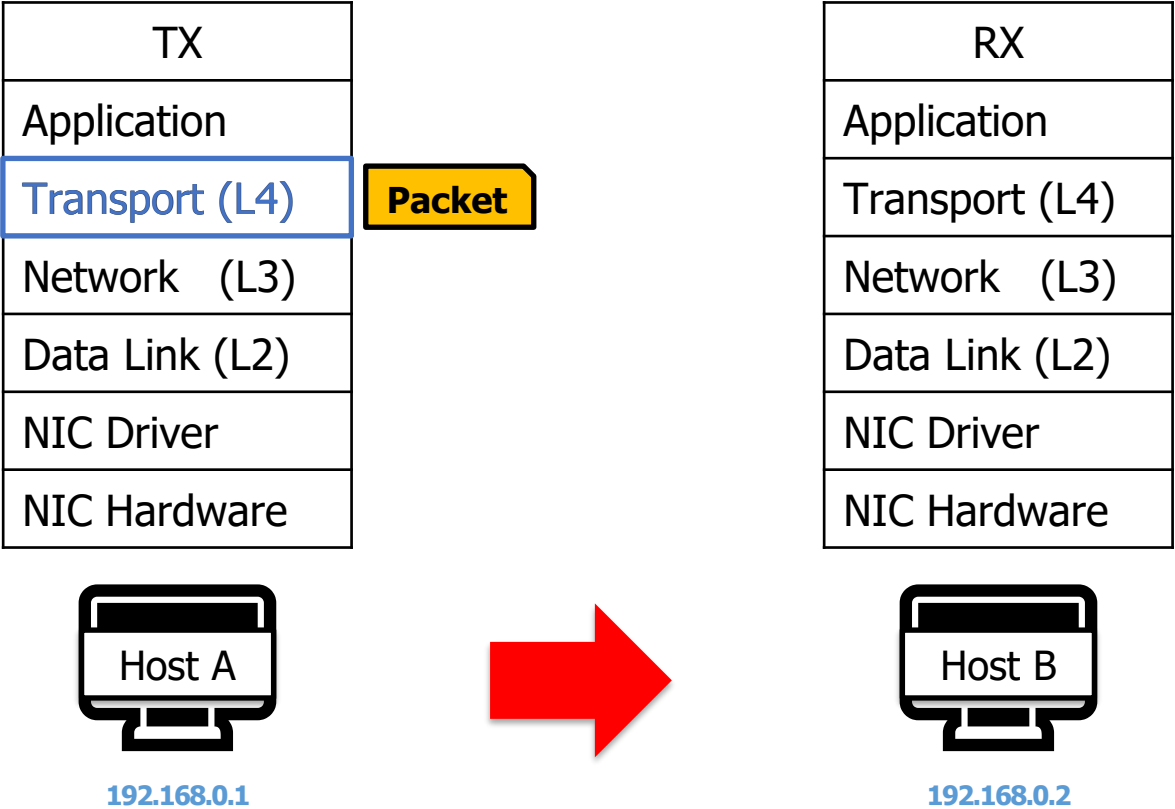
Packet Sending Path



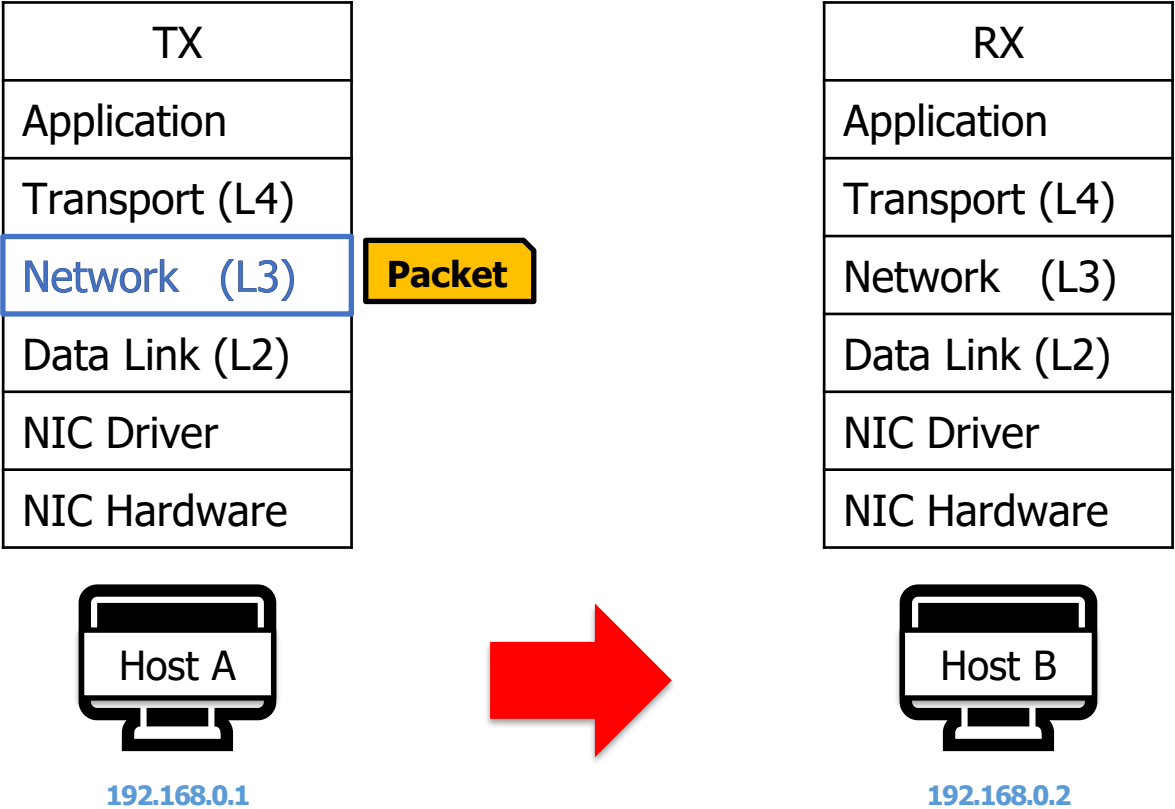
Packet Sending Path



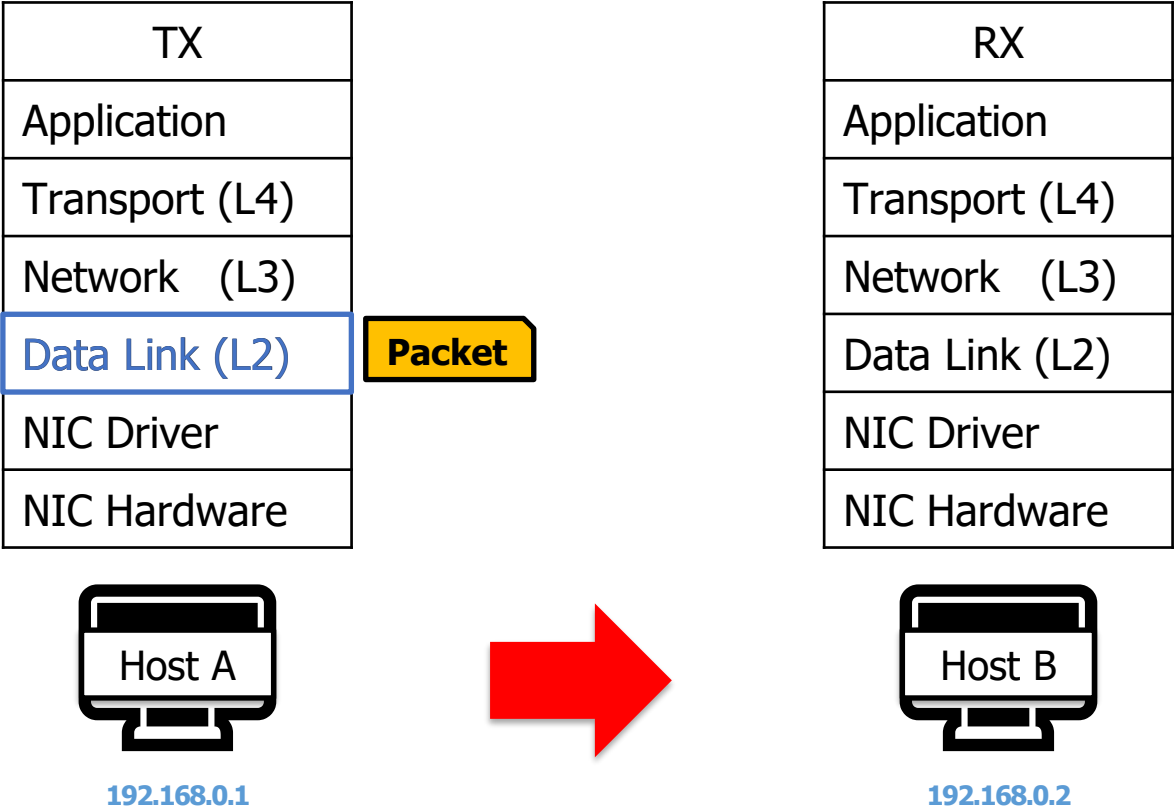
Packet Sending Path



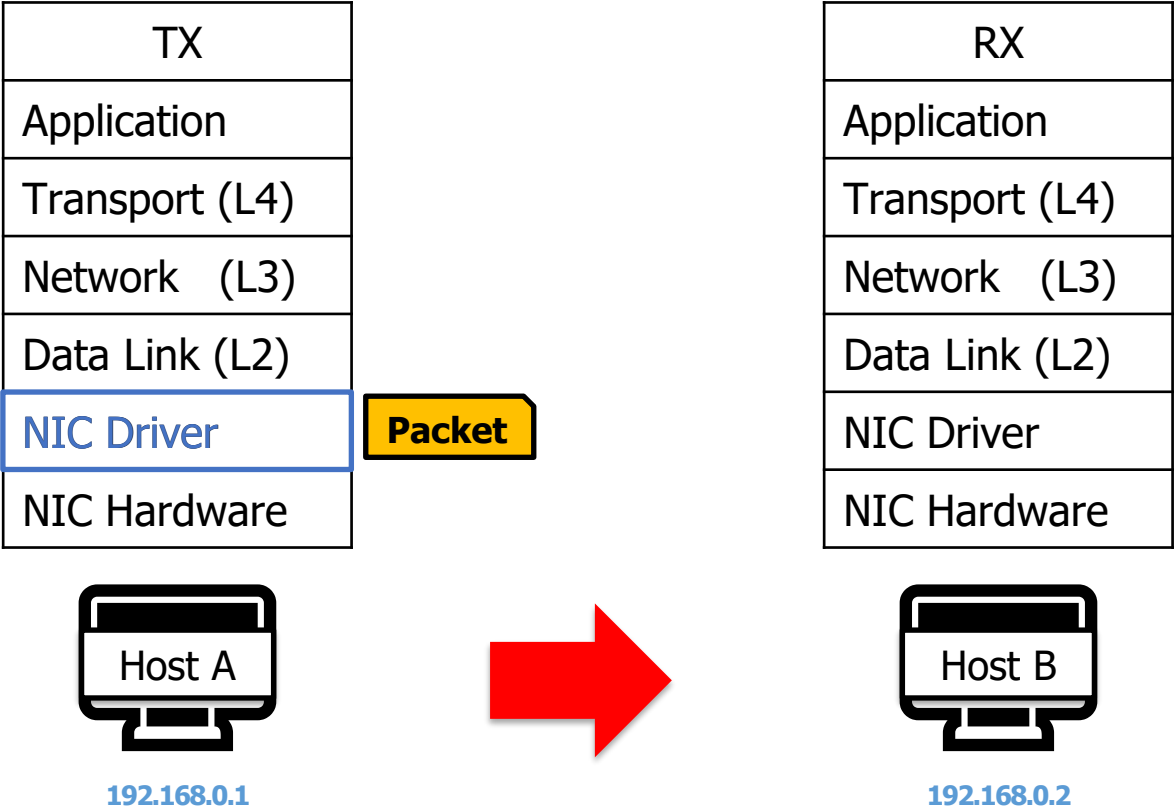
Packet Sending Path



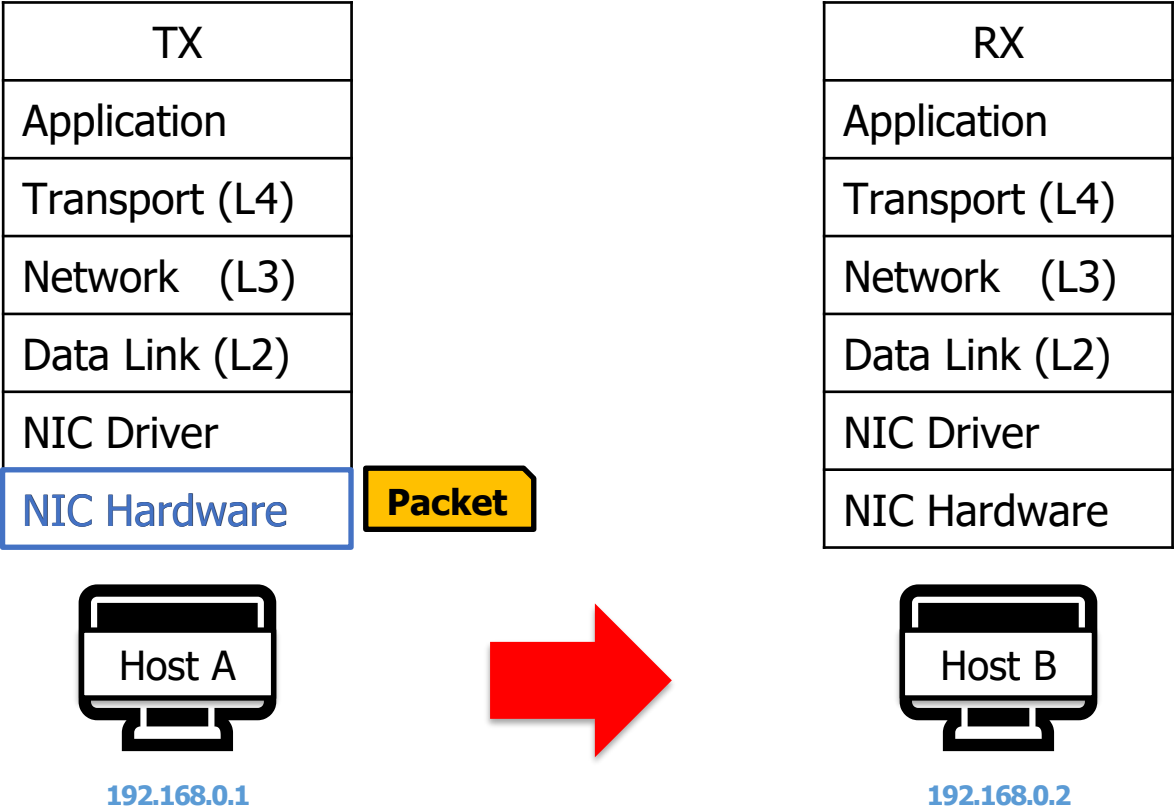
Packet Sending Path



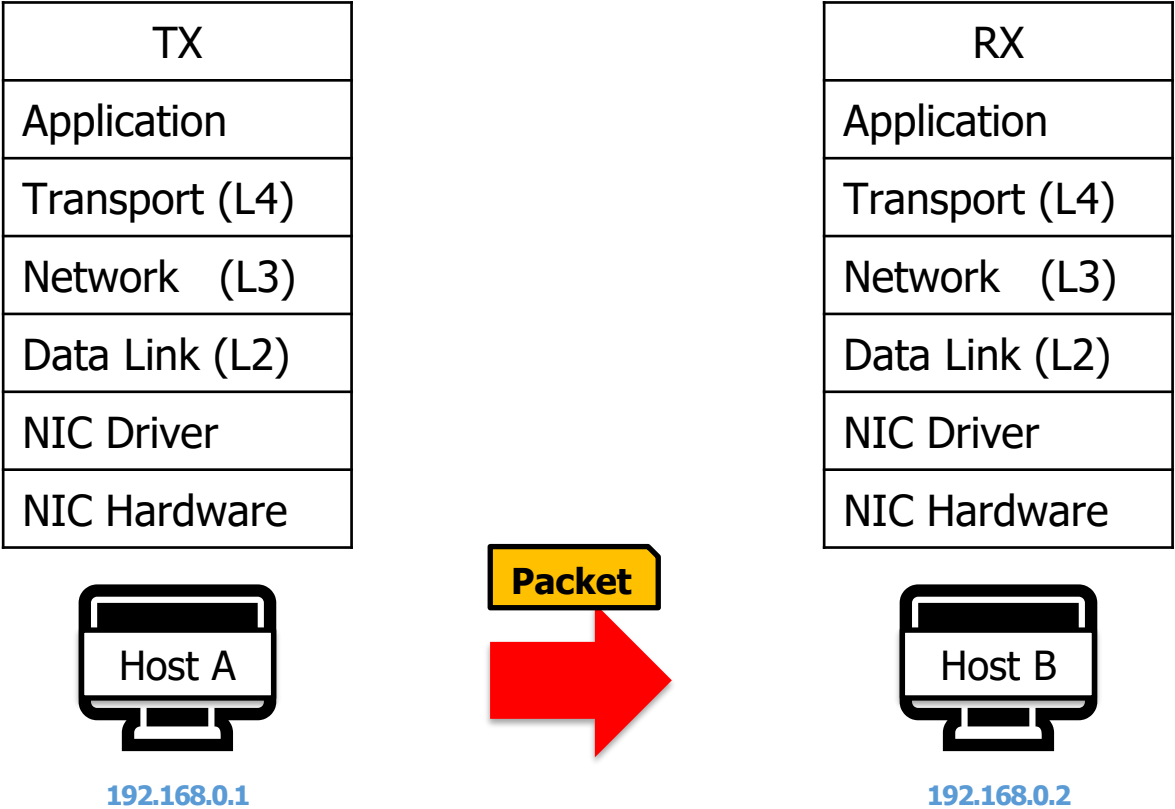
Packet Sending Path



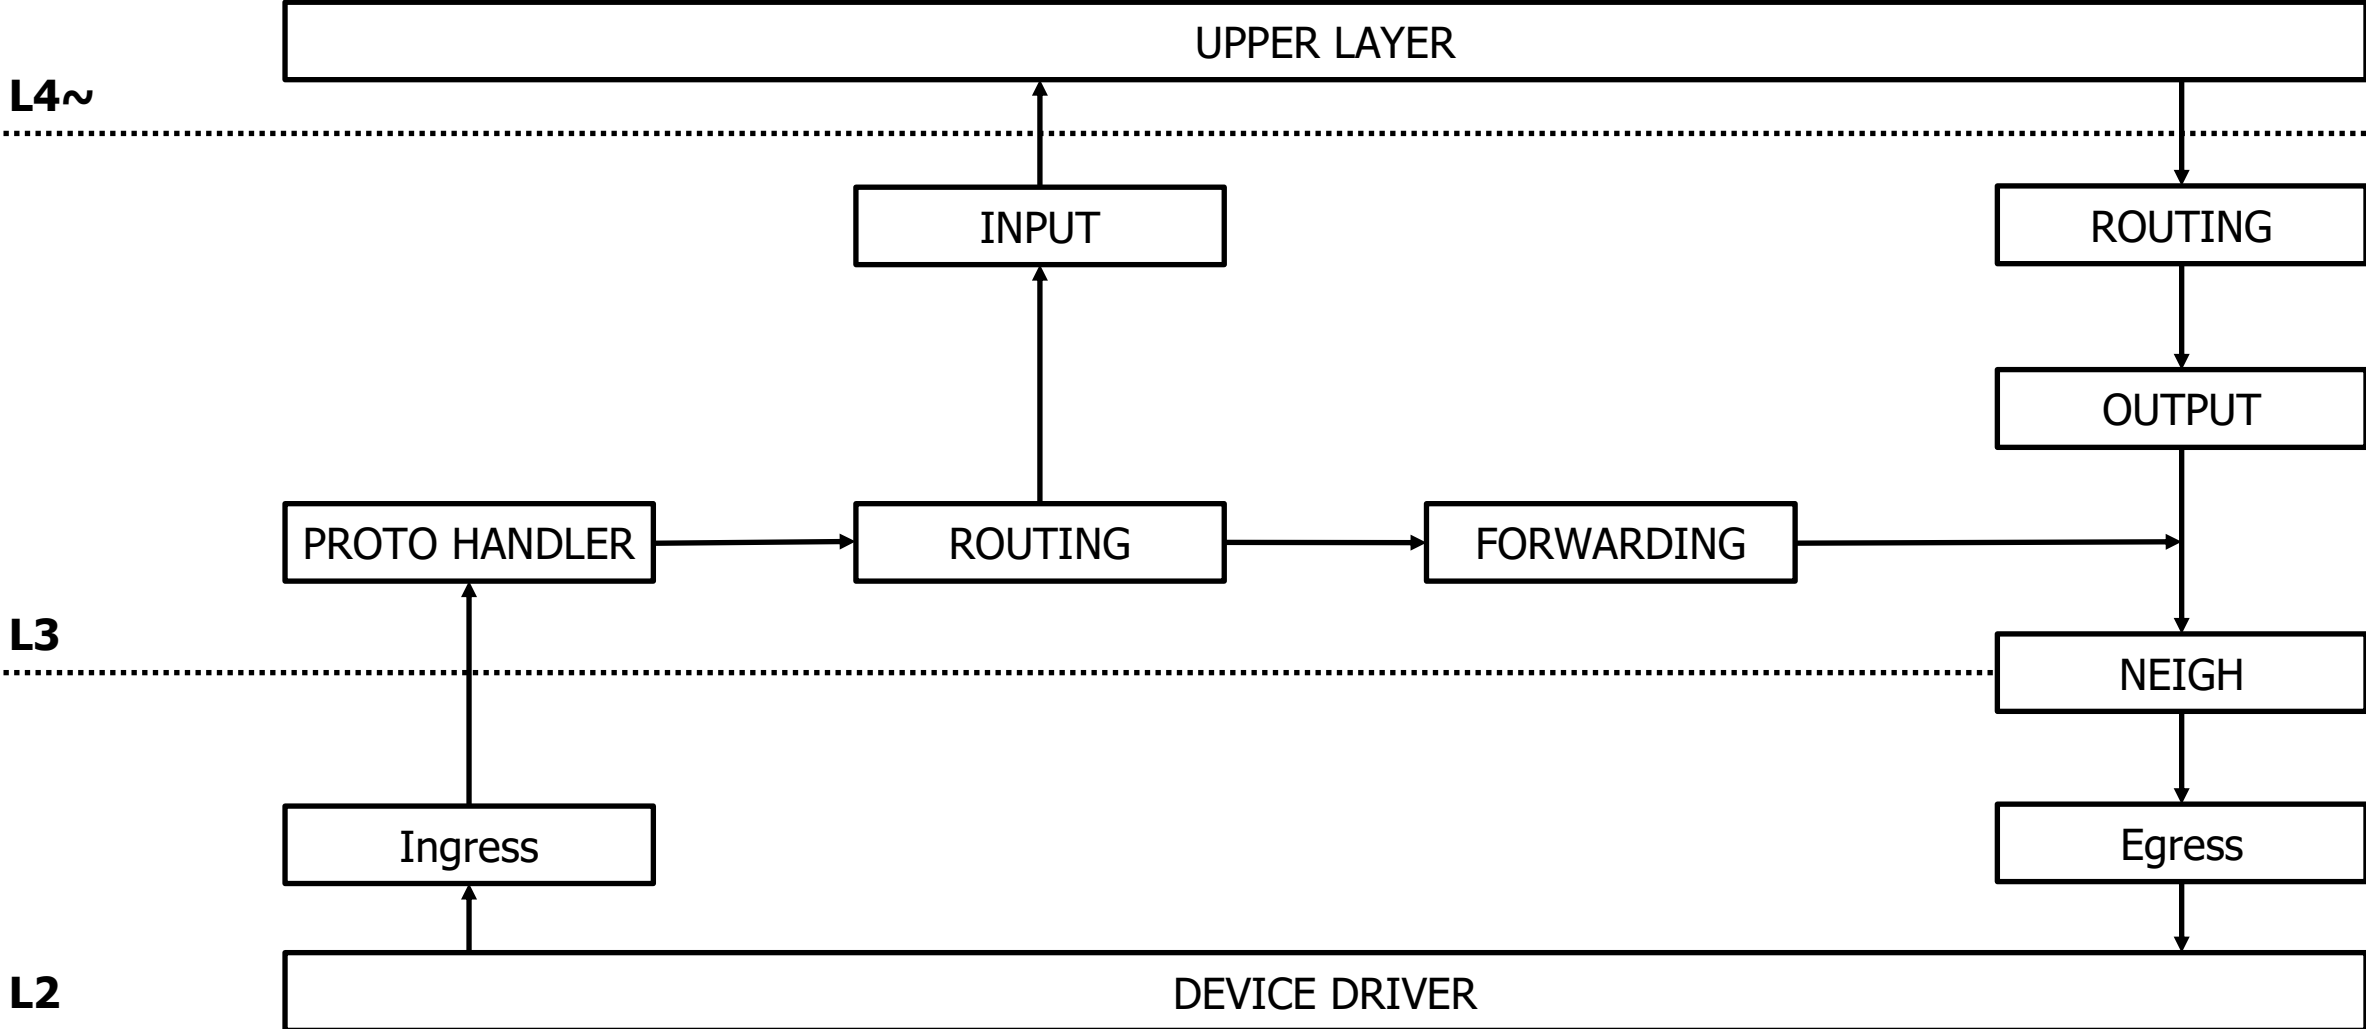
Packet Sending Path




Packet Sending Path



Packet Path in Kernel



Packet Receiving Path



Packet Receiving Path

네트워크 인터페이스 카드 하드웨어에서 이를 수신하고
디바이스 드라이버단에서 NIC의 하드웨어(RX) 큐에서 패킷을 빼 내어서

L4~



L3

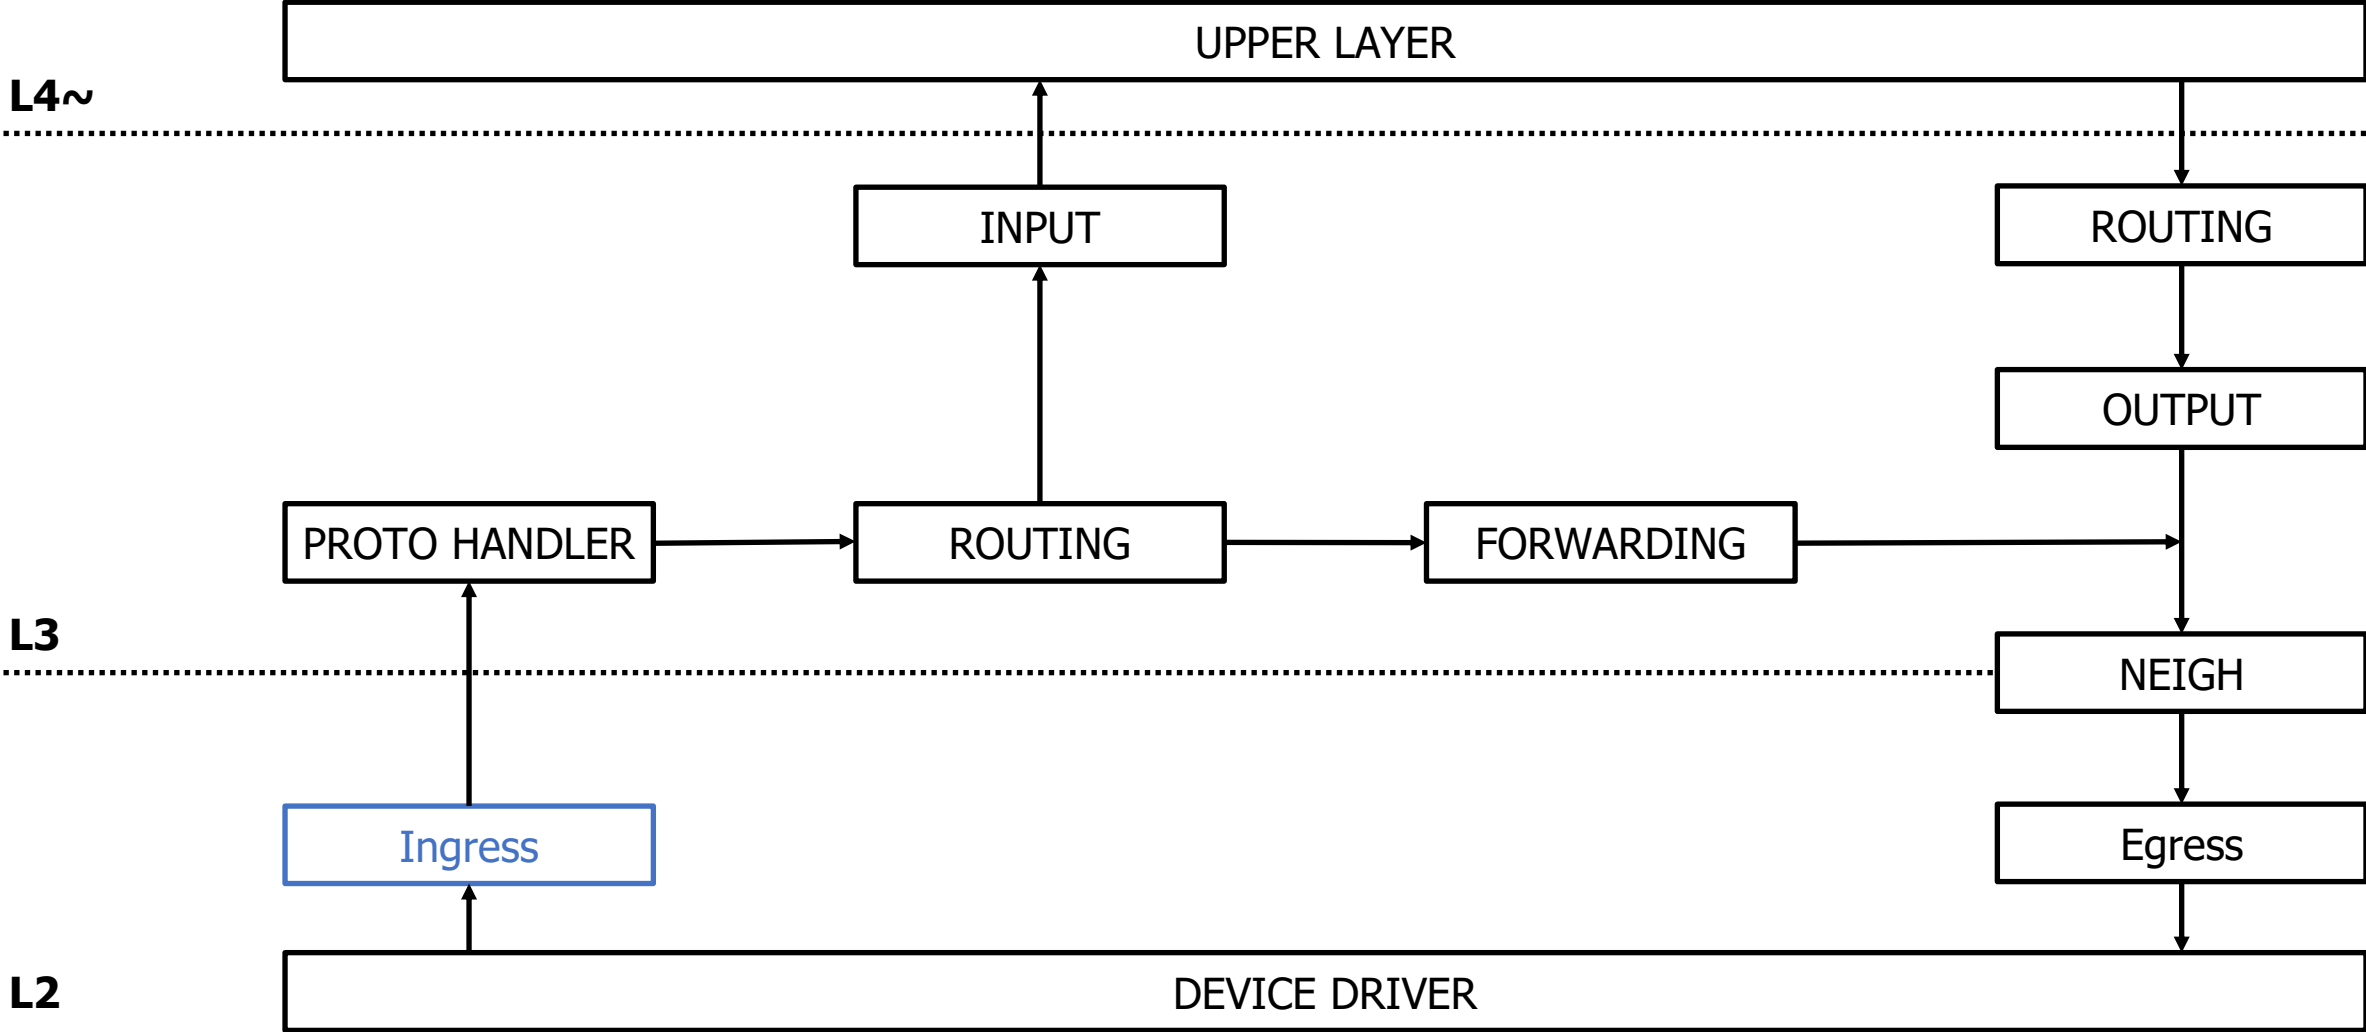


L2



Packet Receiving Path

빠온 패킷을 sk_buff 구조체로 생성하고



Packet Receiving Path

해당 패킷이 어떤 프로토콜인지를 확인하여
해당하는 핸들러 함수를 호출해 프로토콜을 확인

L4~



L3



L2



Packet Receiving Path

ROUTING 테이블에서 나의 IP 주소로 오는 패킷인지 확인

L4~

UPPER LAYER

INPUT

ROUTING

OUTPUT

L3

PROTO HANDLER

ROUTING

FORWARDING

NEIGH

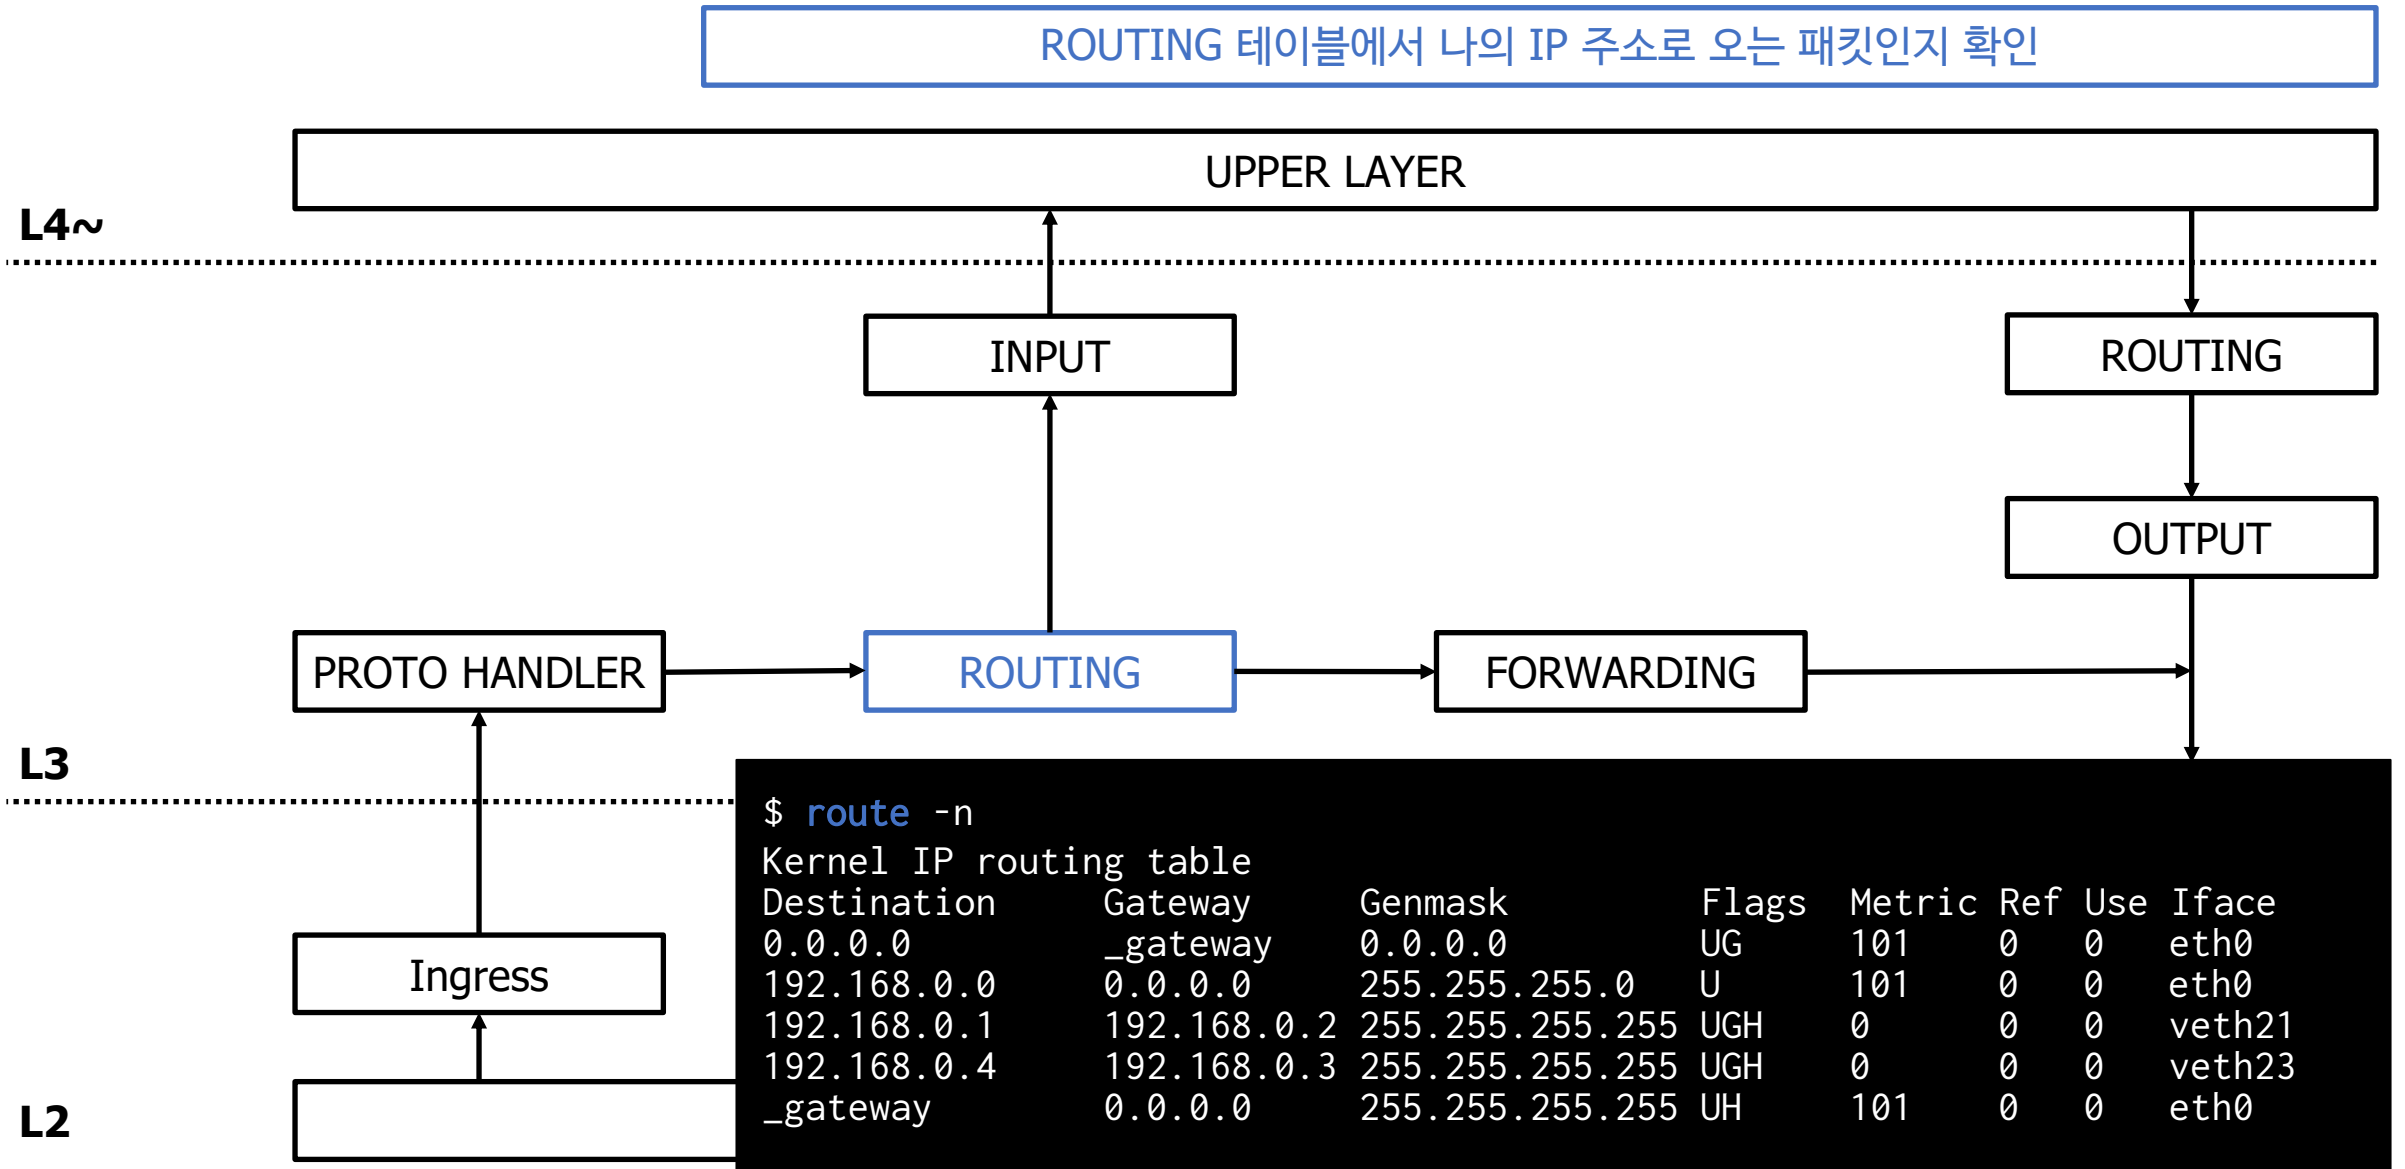
Ingress

Egress

L2

DEVICE DRIVER

Packet Receiving Path



Packet Receiving Path

L4 함수 호출 전 전처리 진행

L4~

UPPER LAYER

INPUT

ROUTING

OUTPUT

L3

PROTO HANDLER

ROUTING

FORWARDING

NEIGH

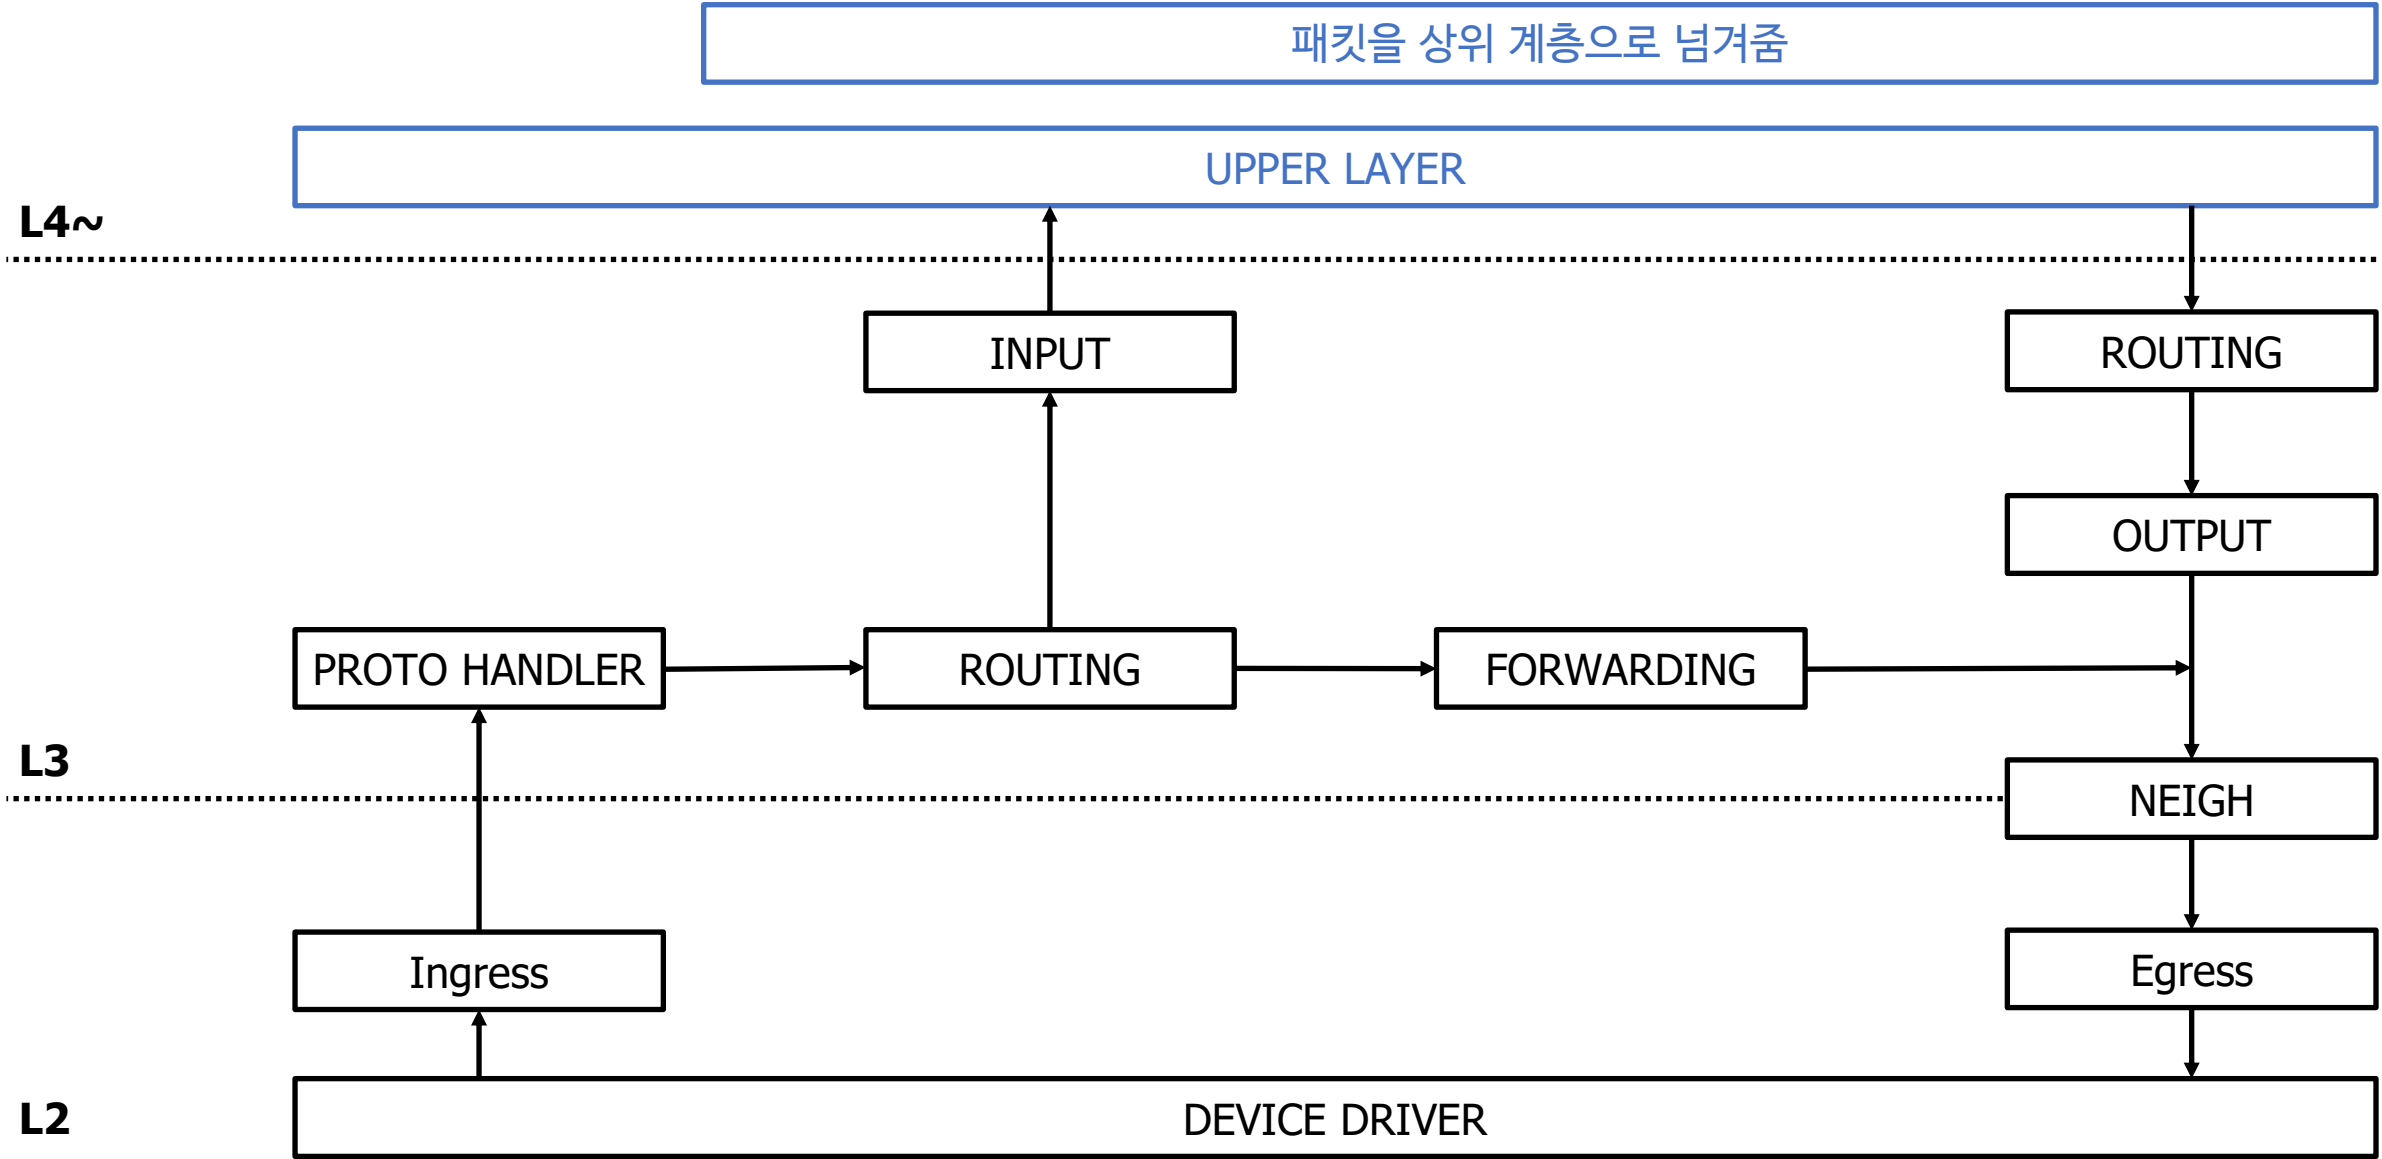
Ingress

Egress

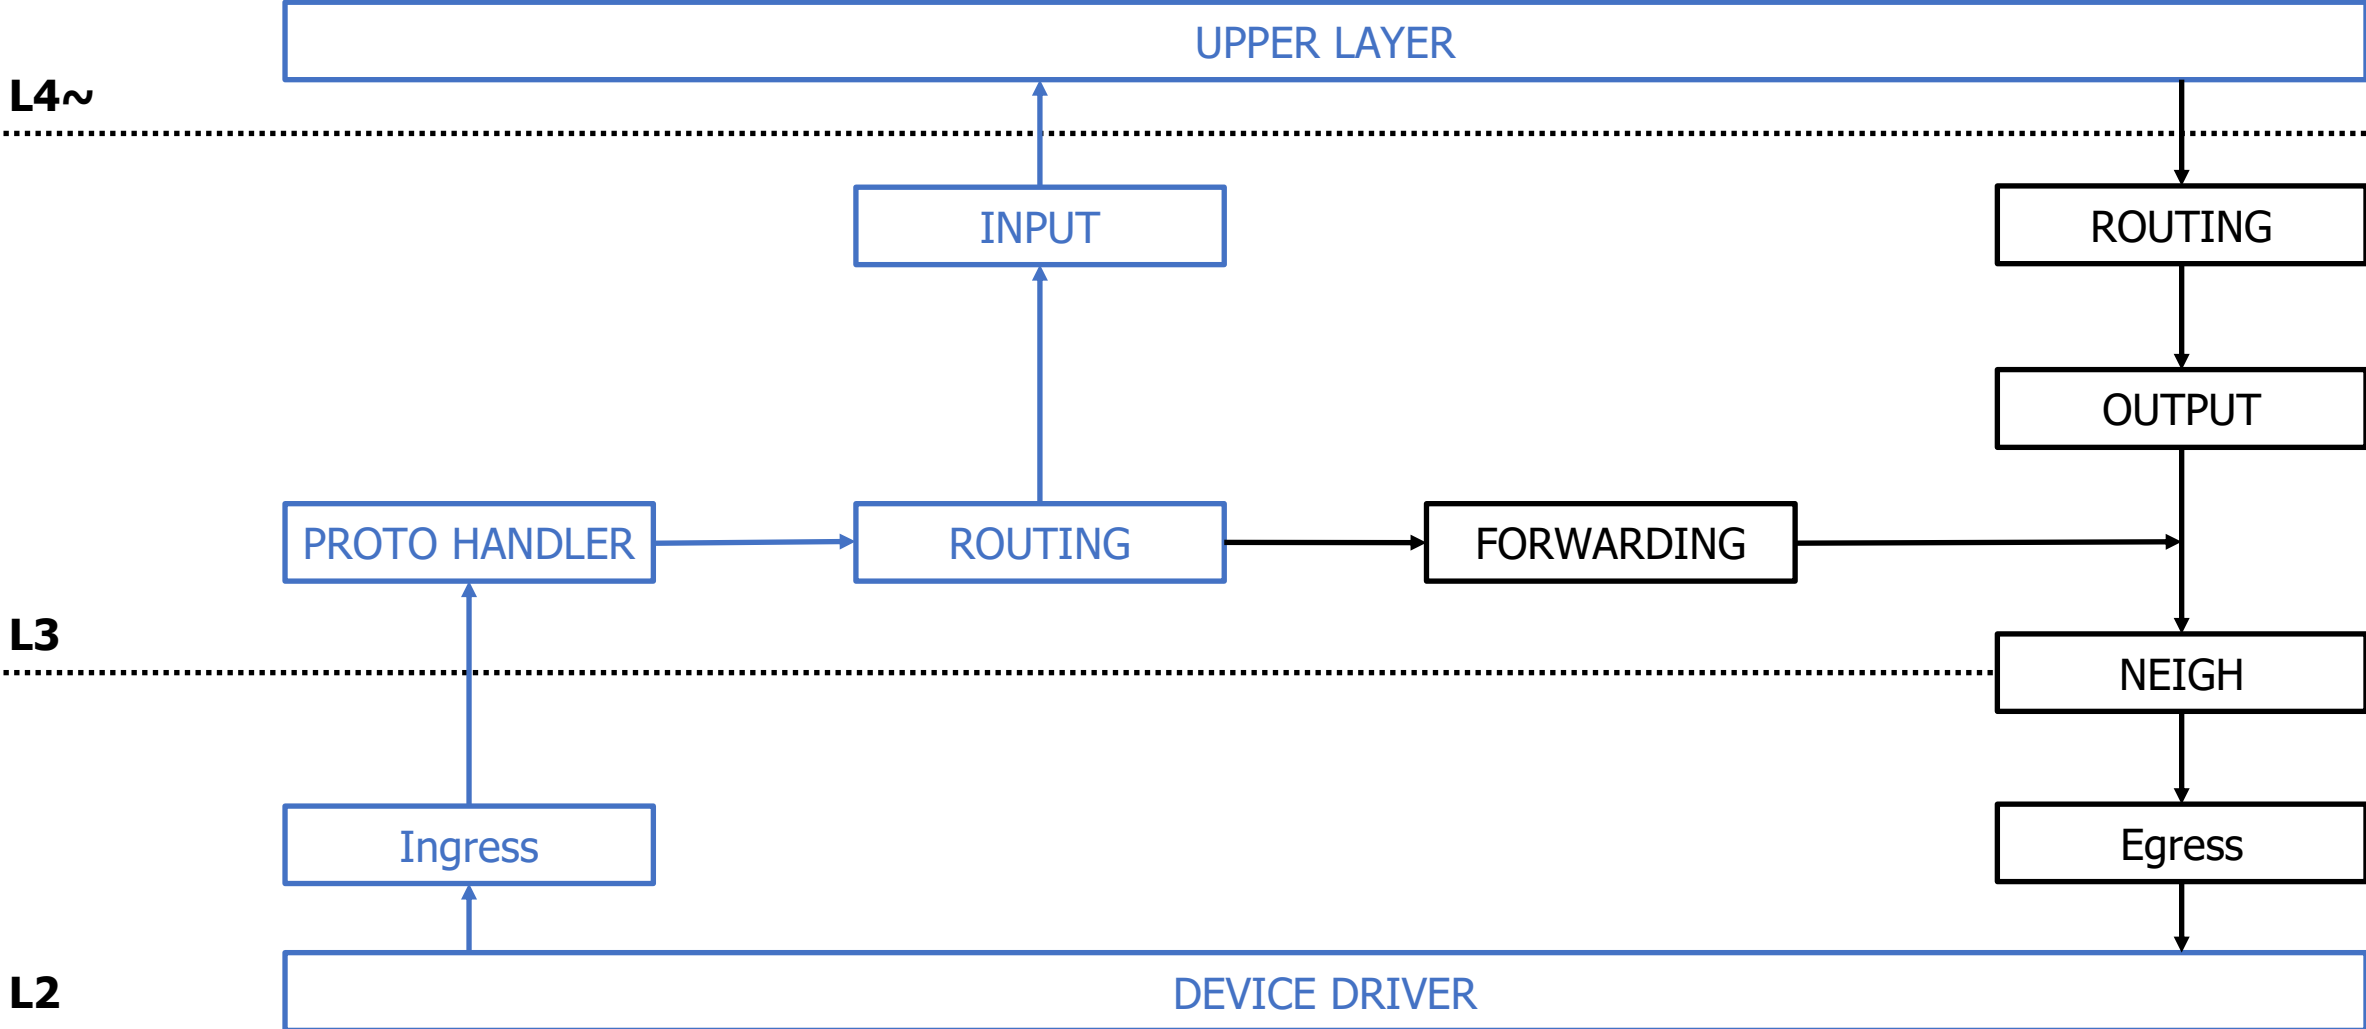
L2

DEVICE DRIVER


Packet Receiving Path



Packet Receiving Path



Packet Forwarding Path



Packet Forwarding Path

ROUTING 테이블에서 나의 IP 주소로 오는 패킷인지 확인

L4~

UPPER LAYER

INPUT

ROUTING

OUTPUT

L3

PROTO HANDLER

ROUTING

FORWARDING

NEIGH

Ingress

Egress

L2

DEVICE DRIVER

Packet Forwarding Path

FORWARDING 인 경우 라우팅 테이블에 따른 인터페이스로 이동

L4~



L3



L2



Packet Forwarding Path

전송할 목적지의 MAC 주소 확인

L4~

UPPER LAYER

INPUT

ROUTING

OUTPUT

L3

PROTO HANDLER

ROUTING

FORWARDING

NEIGH

Ingress

Egress

L2

DEVICE DRIVER

Packet Forwarding Path

sk_buff를 패킷으로 만듦

L4~

UPPER LAYER

INPUT

ROUTING

OUTPUT

L3

PROTO HANDLER

ROUTING

FORWARDING

NEIGH

Ingress

Egress

L2

DEVICE DRIVER

Packet Forwarding Path

네트워크 인터페이스를 통해 패킷을 전송

L4~

UPPER LAYER

INPUT

ROUTING

OUTPUT

L3

PROTO HANDLER

ROUTING

FORWARDING

NEIGH

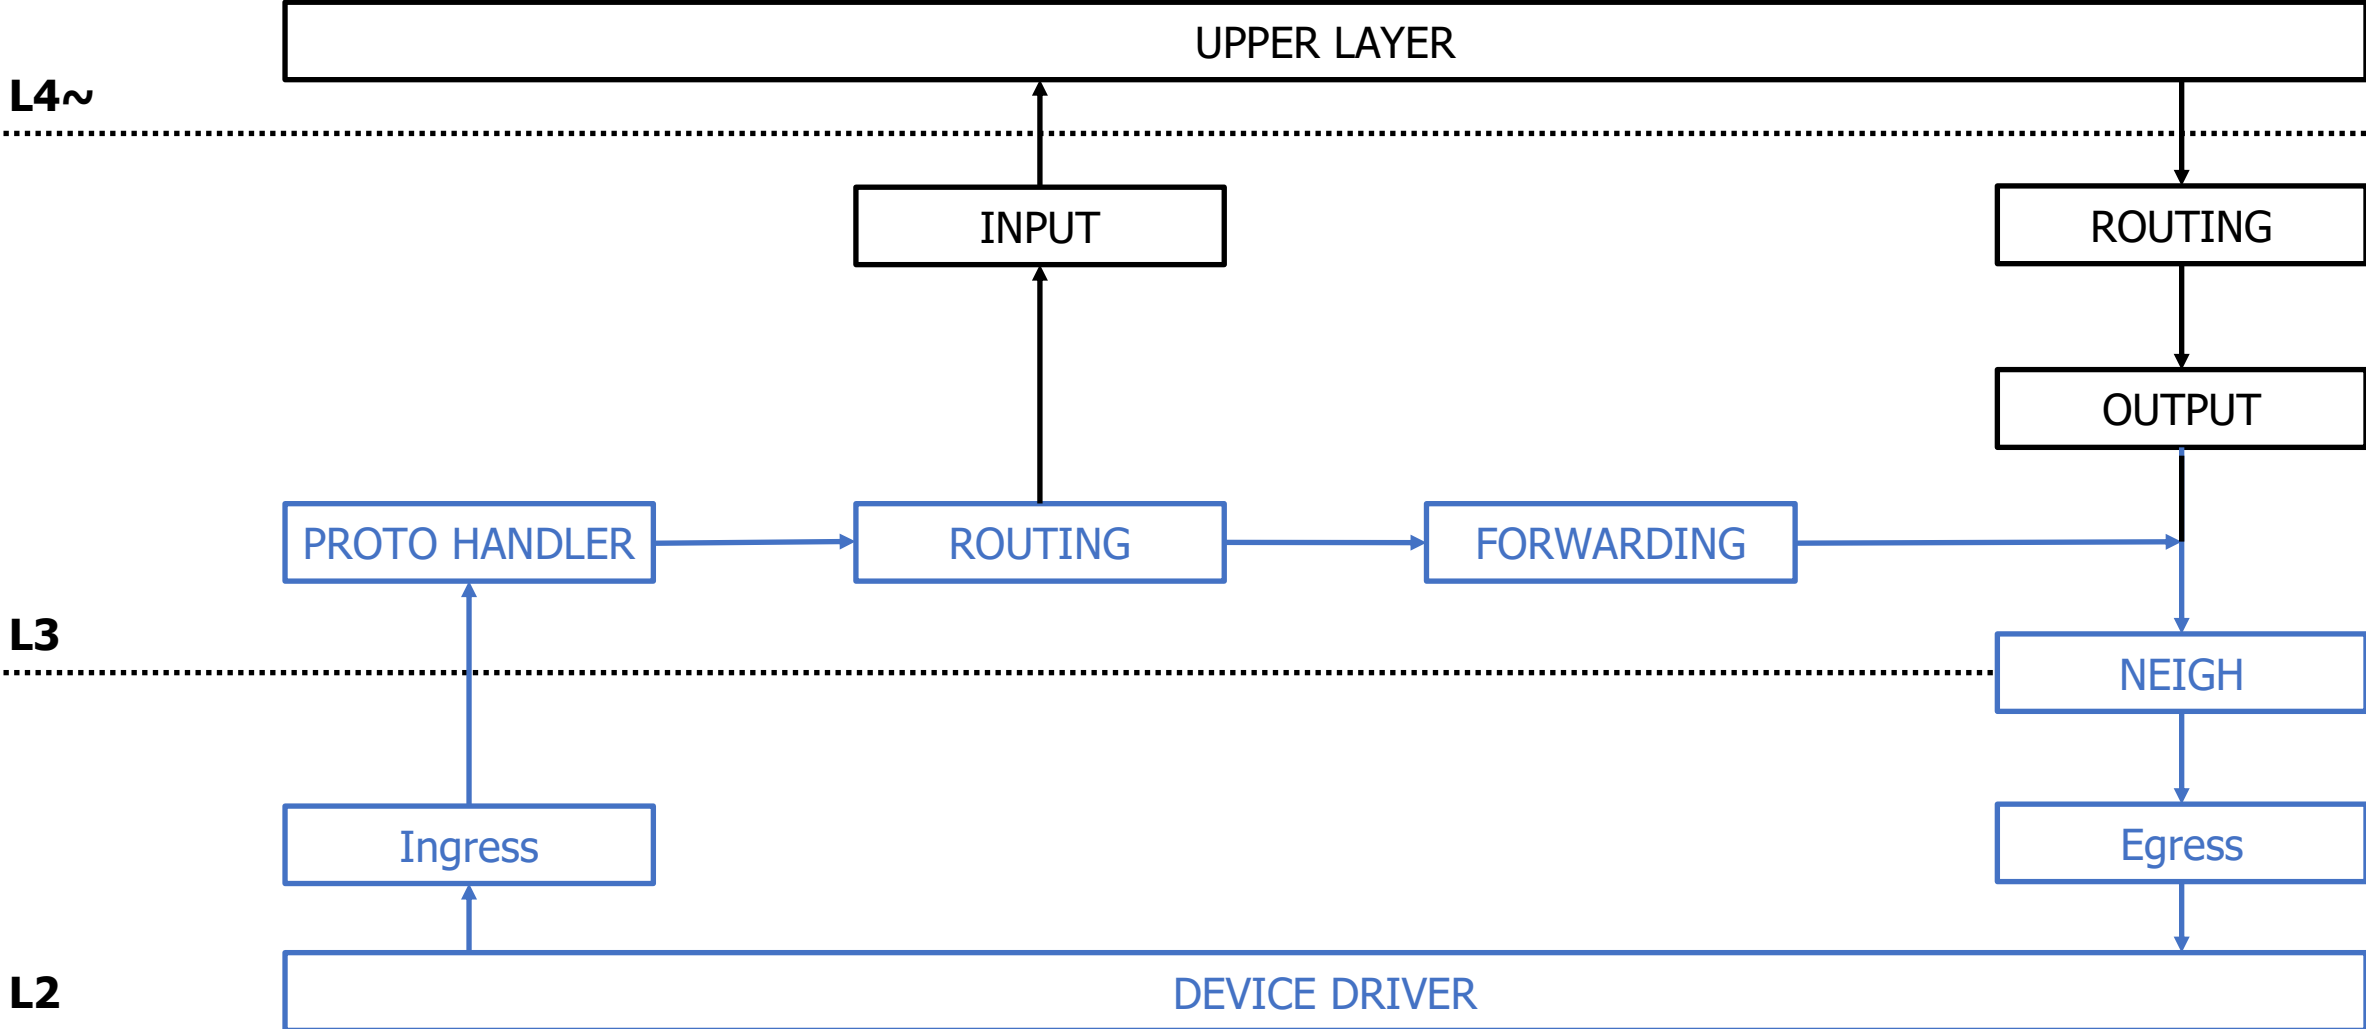
Ingress

Egress


L2

DEVICE DRIVER

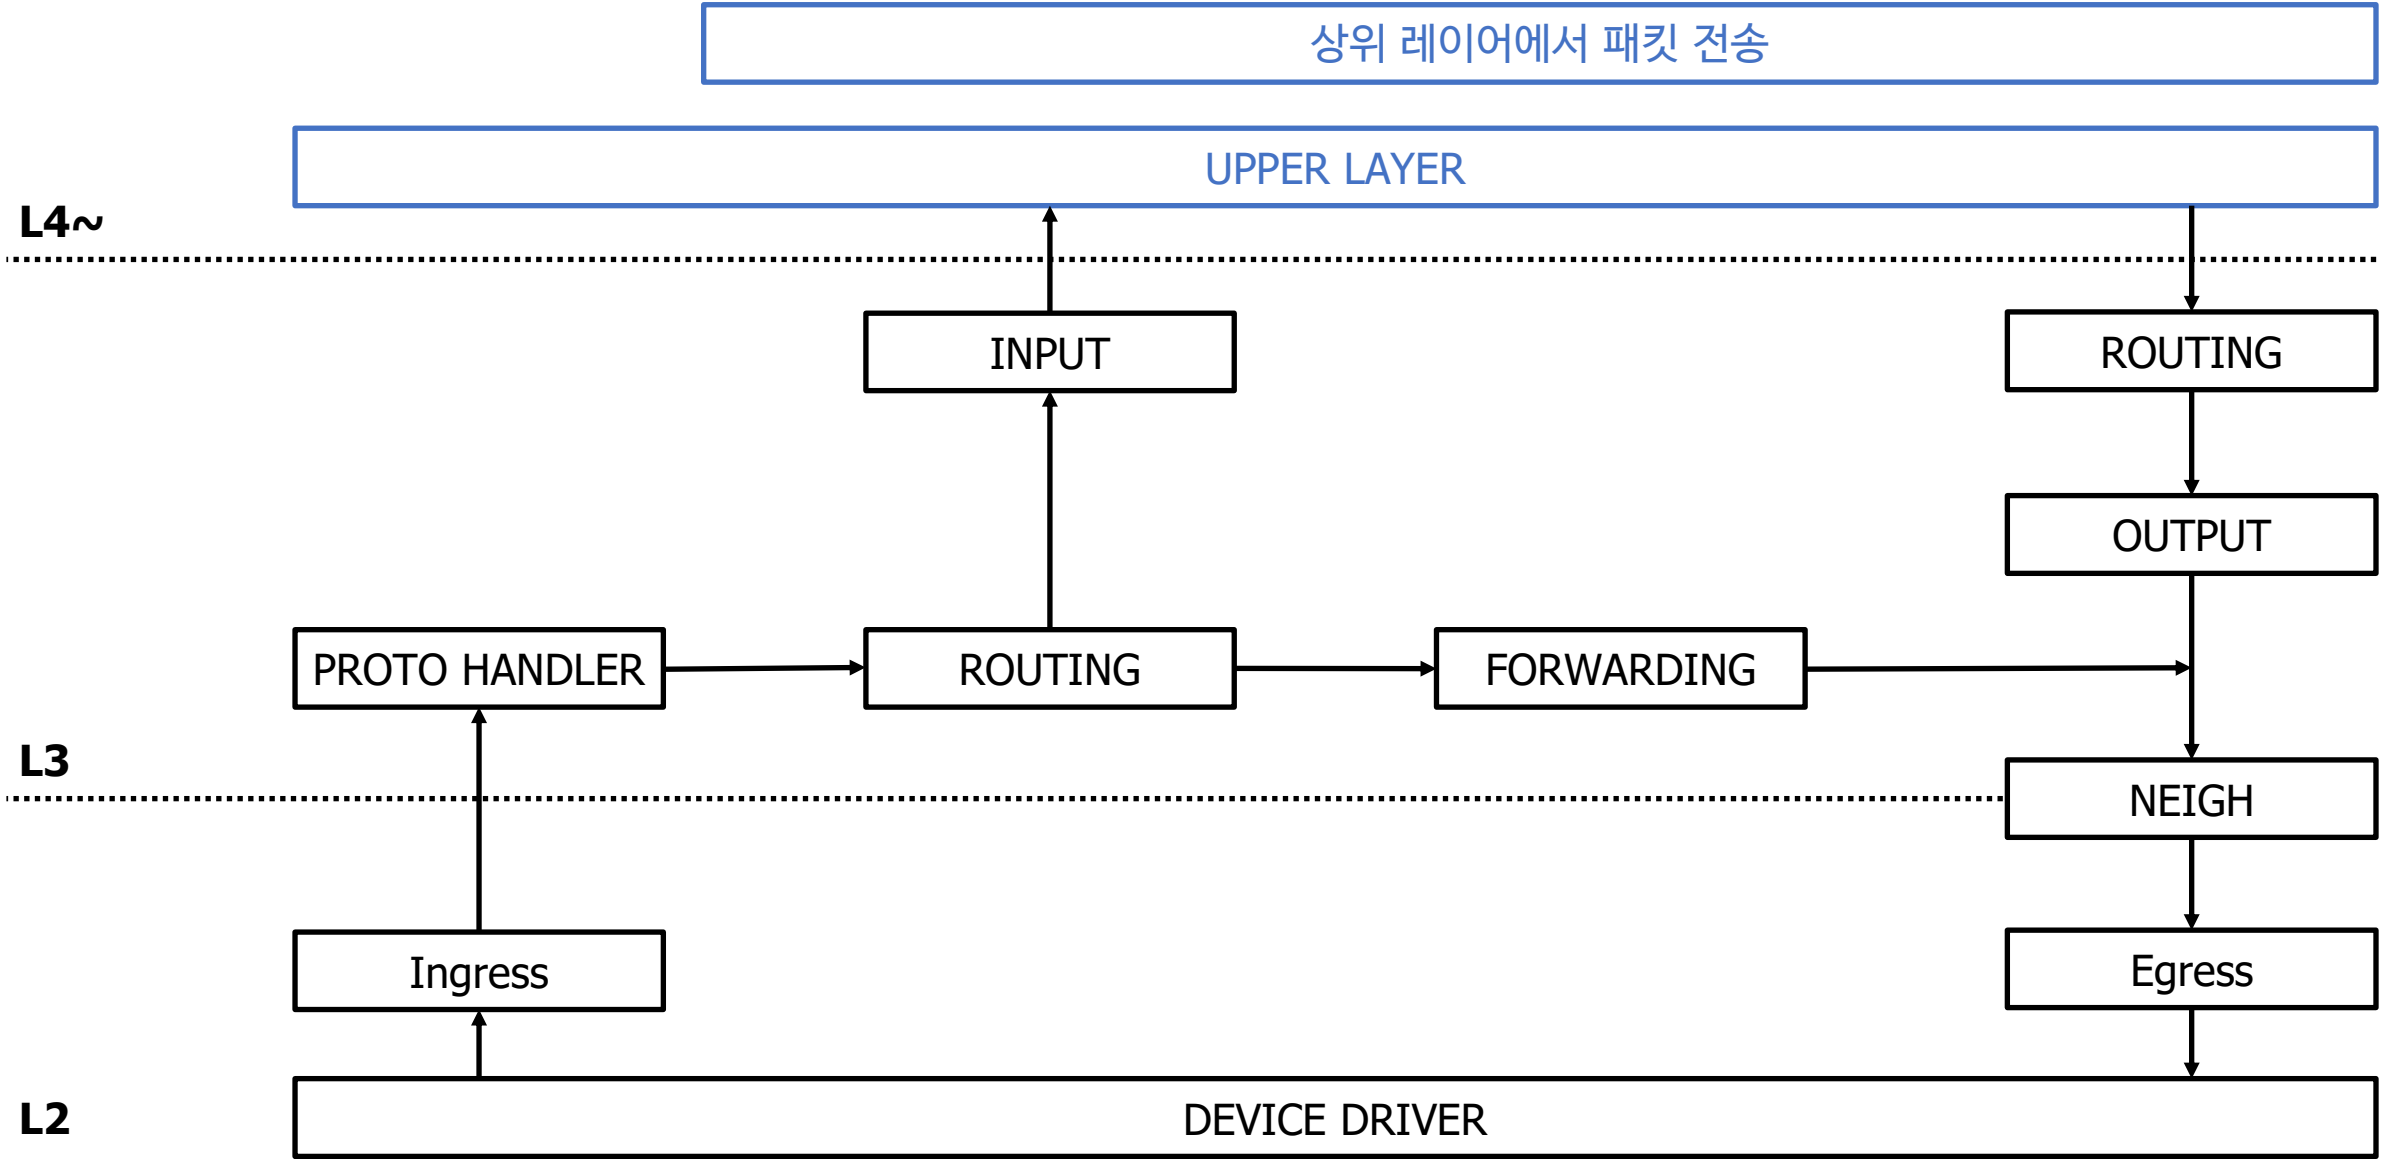
Packet Forwarding Path



Packet Sending Path



Packet Sending Path



Packet Sending Path

ROUTING TABLE을 확인해서 어떤 인터페이스로 전송해야 되는지 확인

L4~



L3



L2



Packet Sending Path

패킷 전송을 위한 전처리 진행

L4~

UPPER LAYER

INPUT

ROUTING

OUTPUT

L3

PROTO HANDLER

ROUTING

FORWARDING

NEIGH

Ingress

Egress

L2

DEVICE DRIVER

Packet Sending Path

전송할 목적지의 MAC 주소 확인

L4~

UPPER LAYER

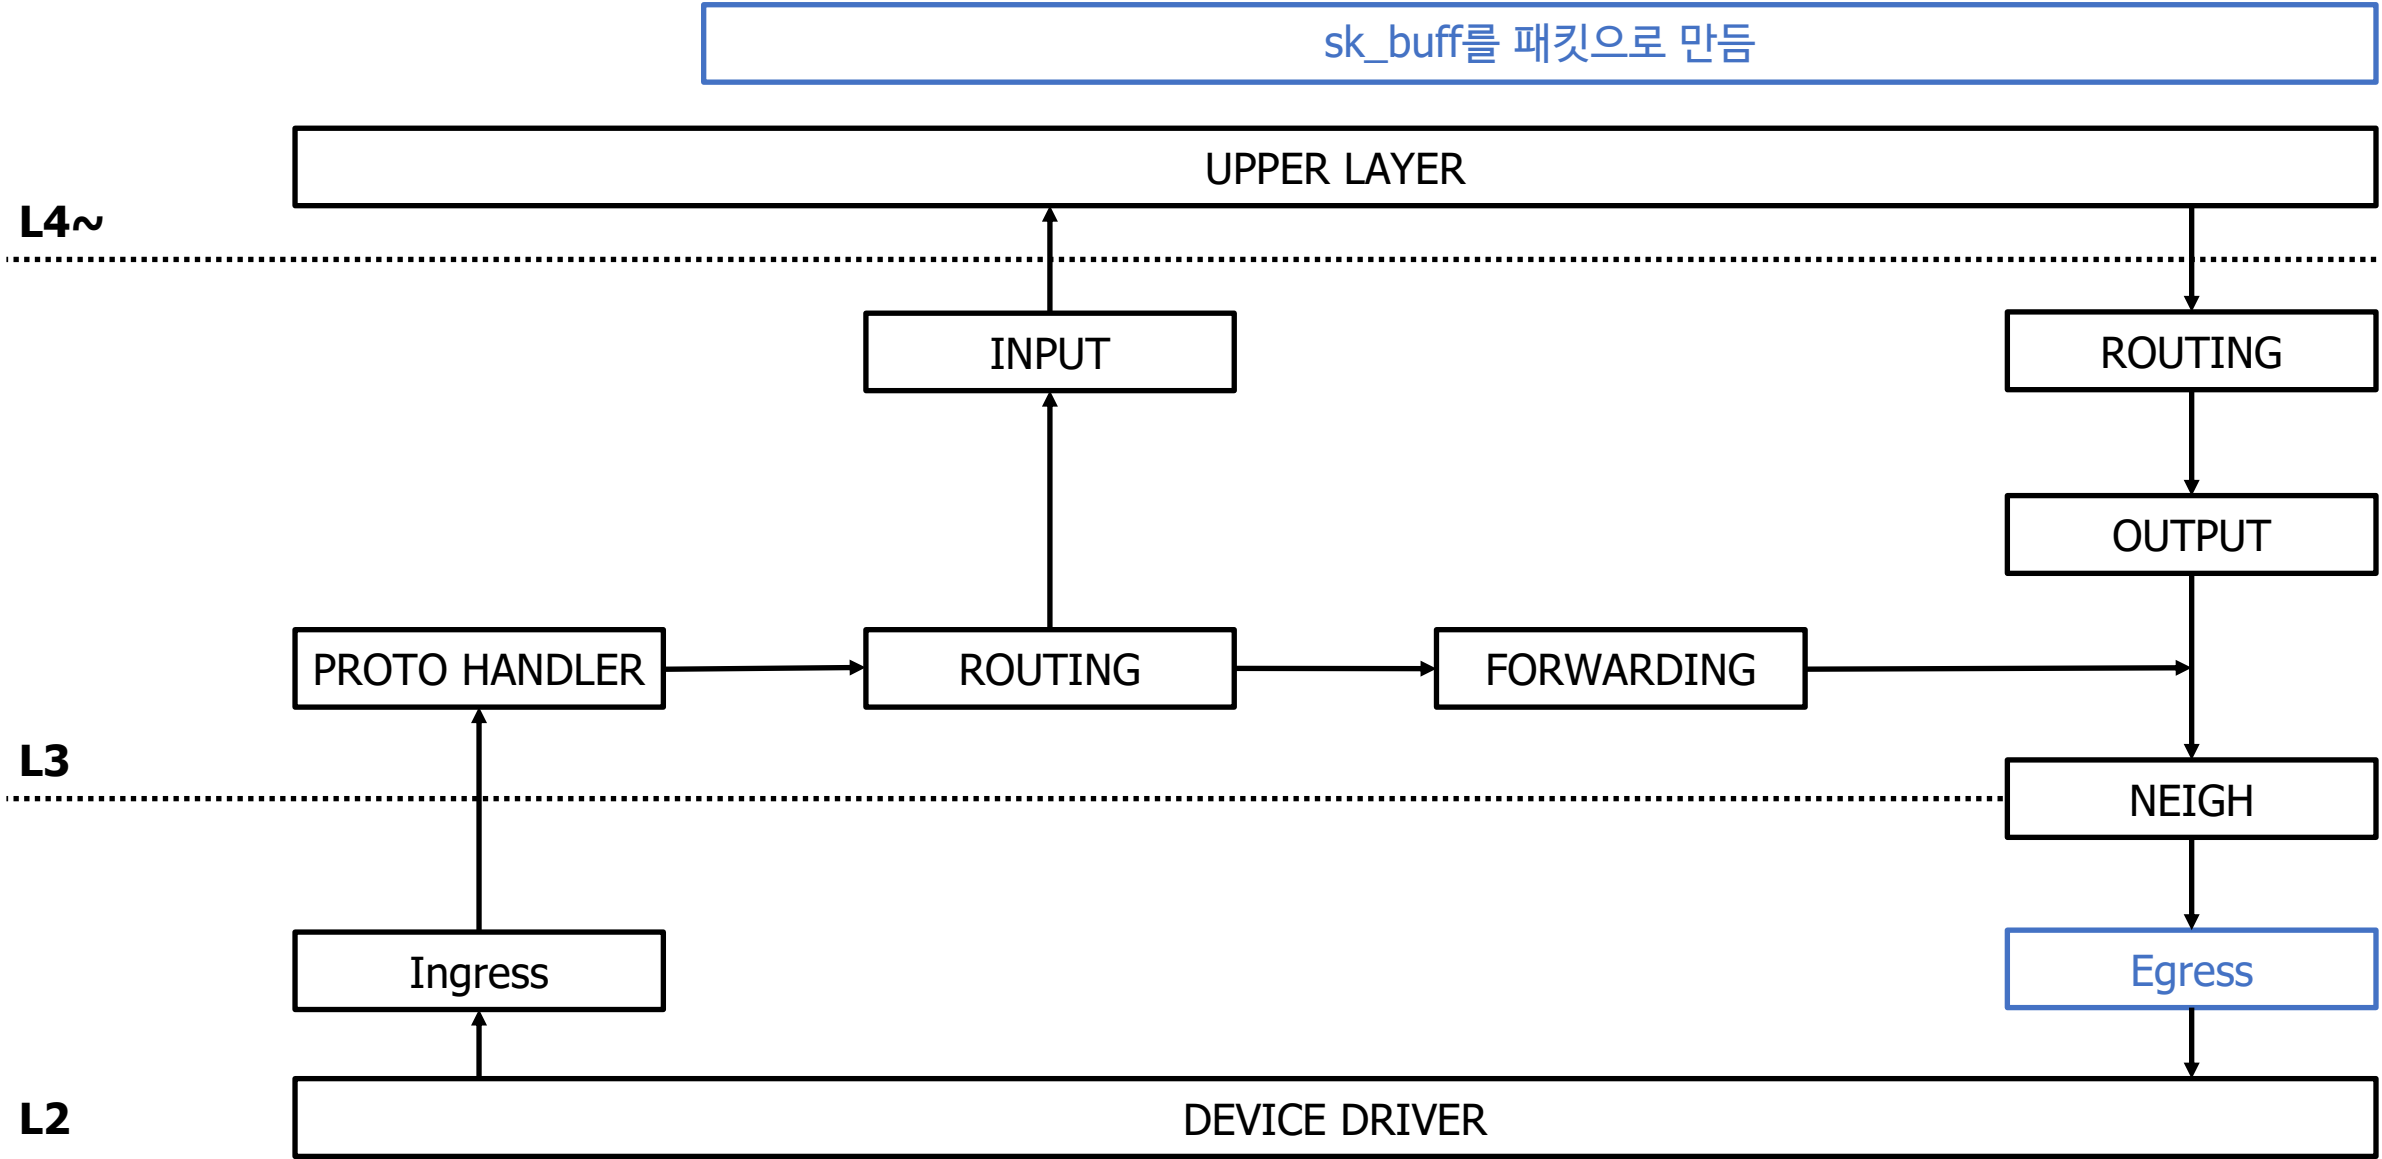
L3

PROTO HANDLER ROUTING FORWARDING NEIGH OUTPUT

L2

Ingress Egress DEVICE DRIVER

Packet Sending Path



Packet Sending Path

네트워크 인터페이스를 통해 패킷을 전송

L4~

UPPER LAYER

INPUT

ROUTING

OUTPUT

L3

PROTO HANDLER

ROUTING

FORWARDING

NEIGH

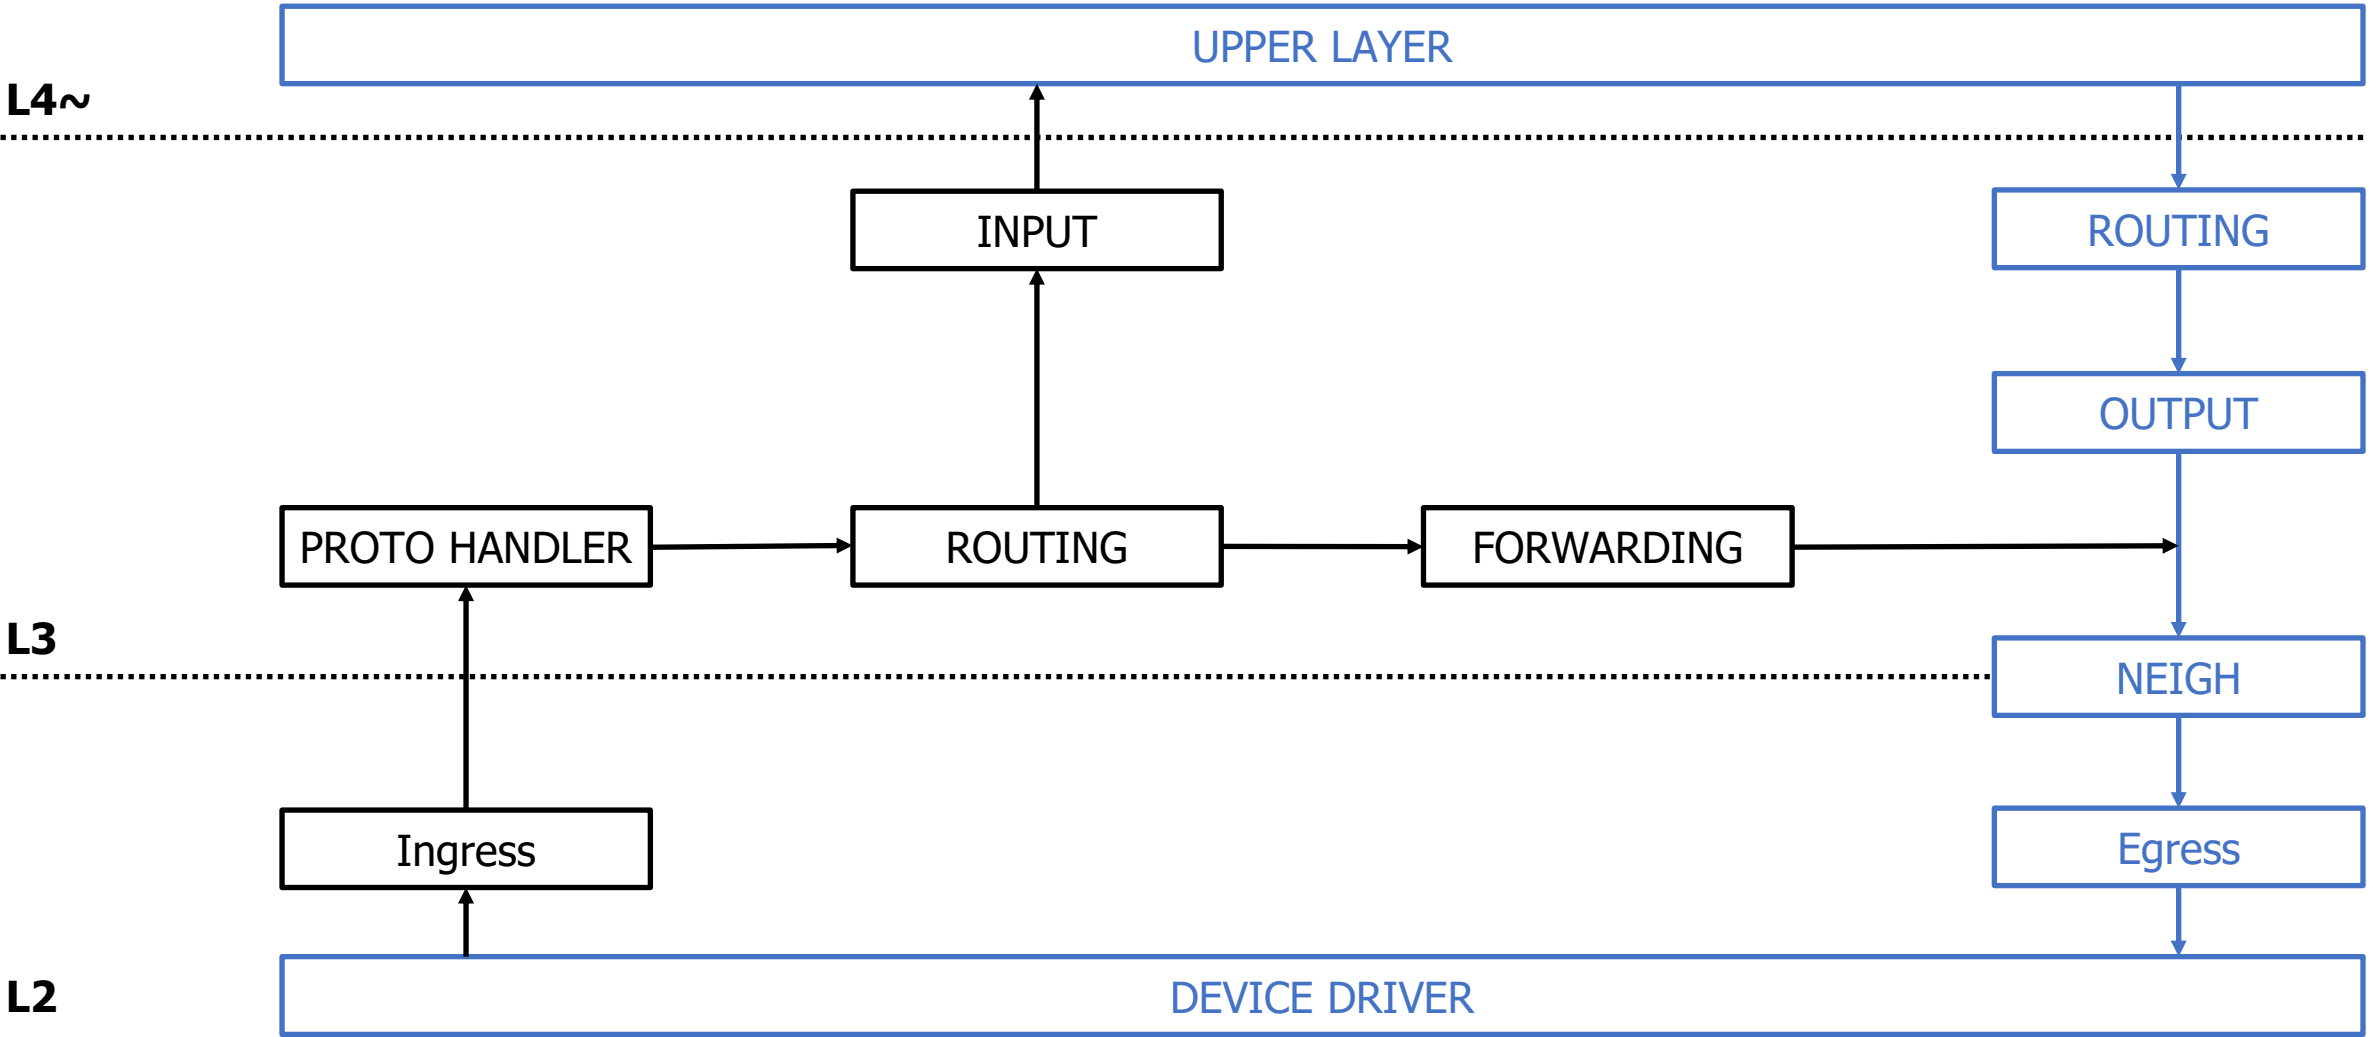
Ingress

Egress

L2


DEVICE DRIVER

Packet Sending Path




Linux Kernel Networking Stack Packet Processing


① Receiving Path



② Forwarding Path



③ Sending Path




Packet Sending

Packet Receiving

Linux Kernel Networking Stack

Packet Forwarding

Linux Kernel Networking Stack



NAT

Packet Sending

Routing

HSR

TCP

Packet Receiving

Packet Processing

Packet Filtering

AMT

Packet Forwarding

UDP

Packet Mangling

Packet Tunneling

NAT

Packet Sending

Routing

HSR

TCP

Packet Receiving

Packet Processing

Packet Filtering

AMT

Packet Forwarding

UDP

Packet Mangling

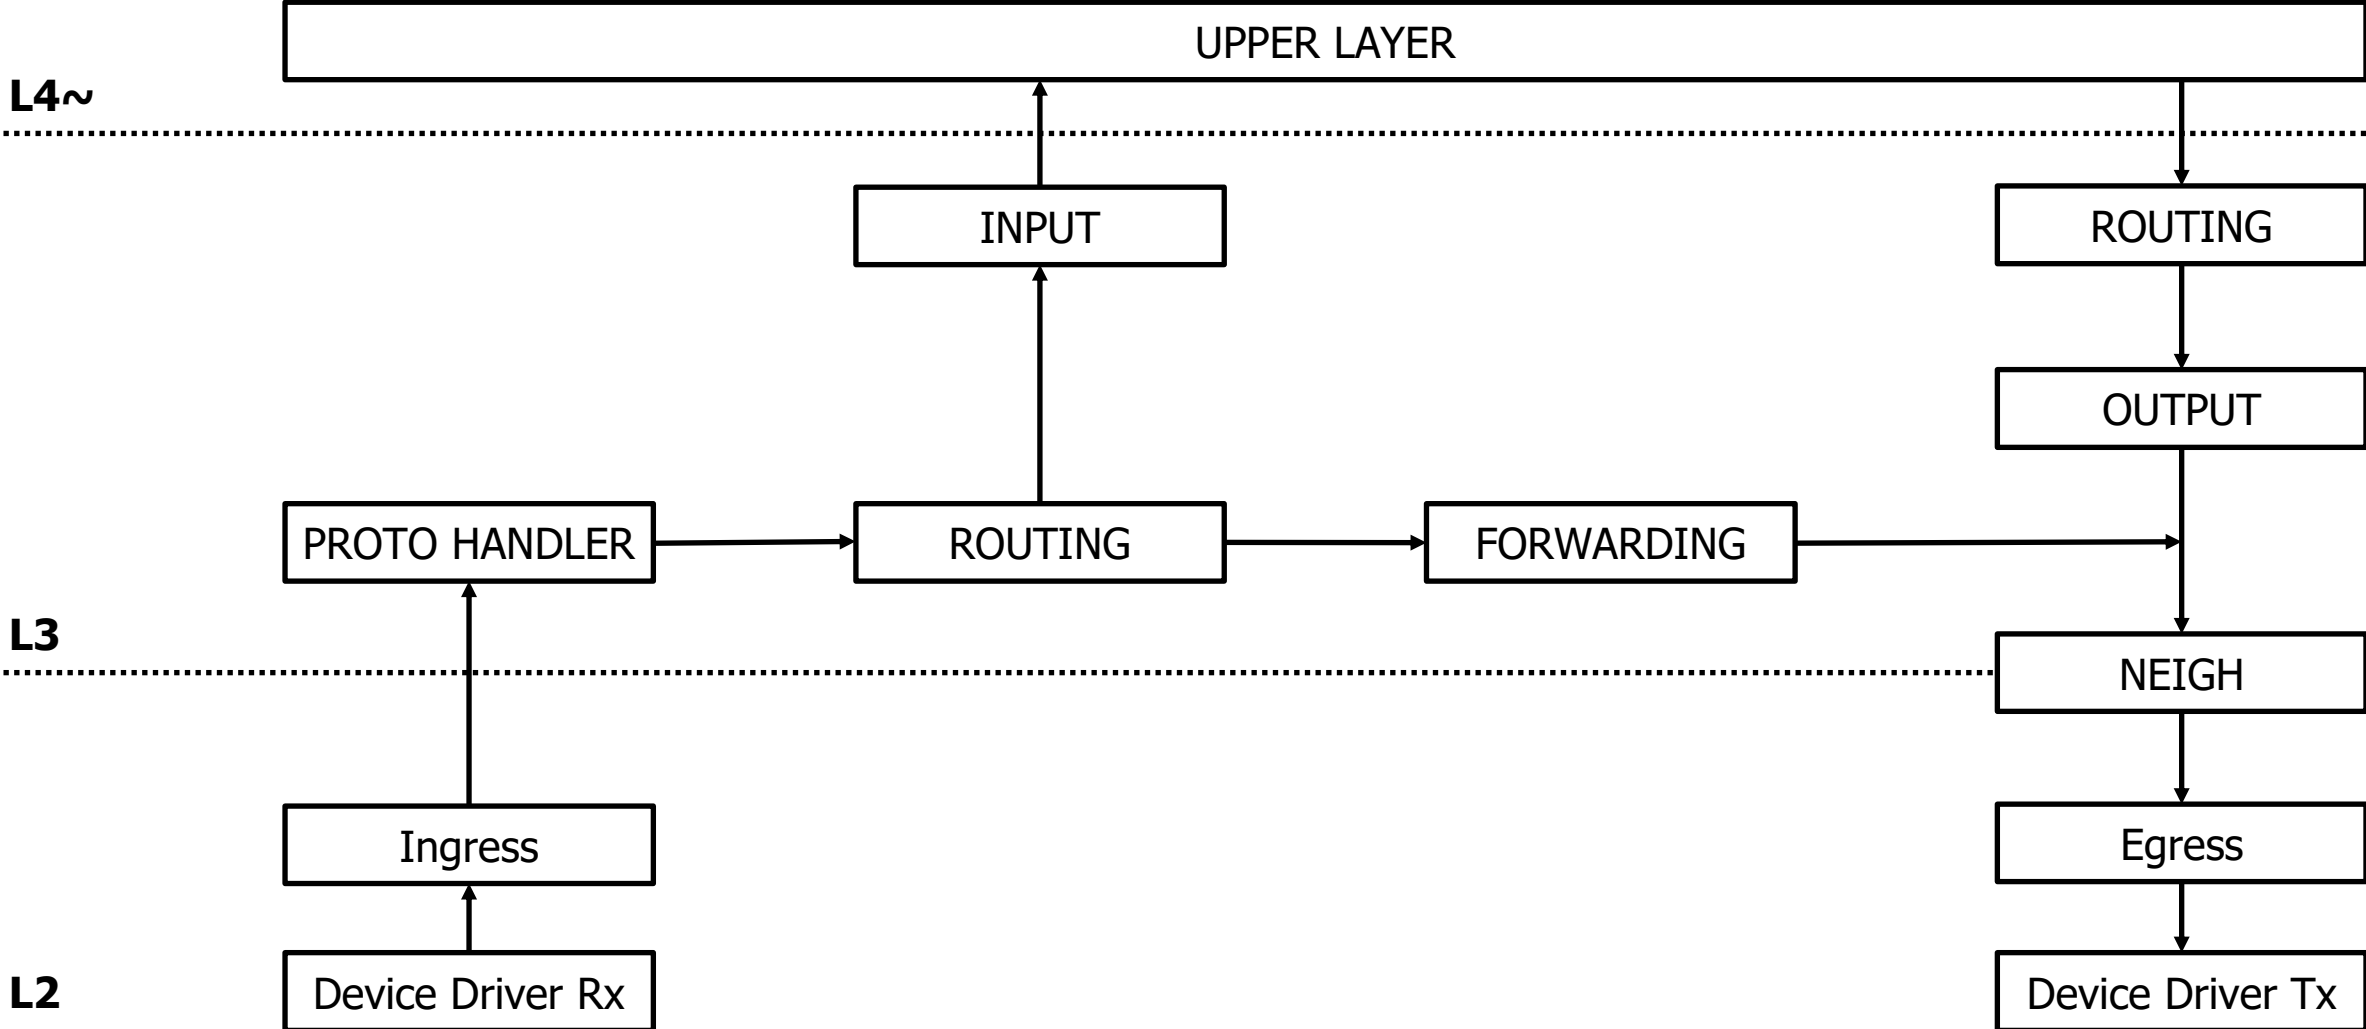
Packet Tunneling

Packet Filtering Path

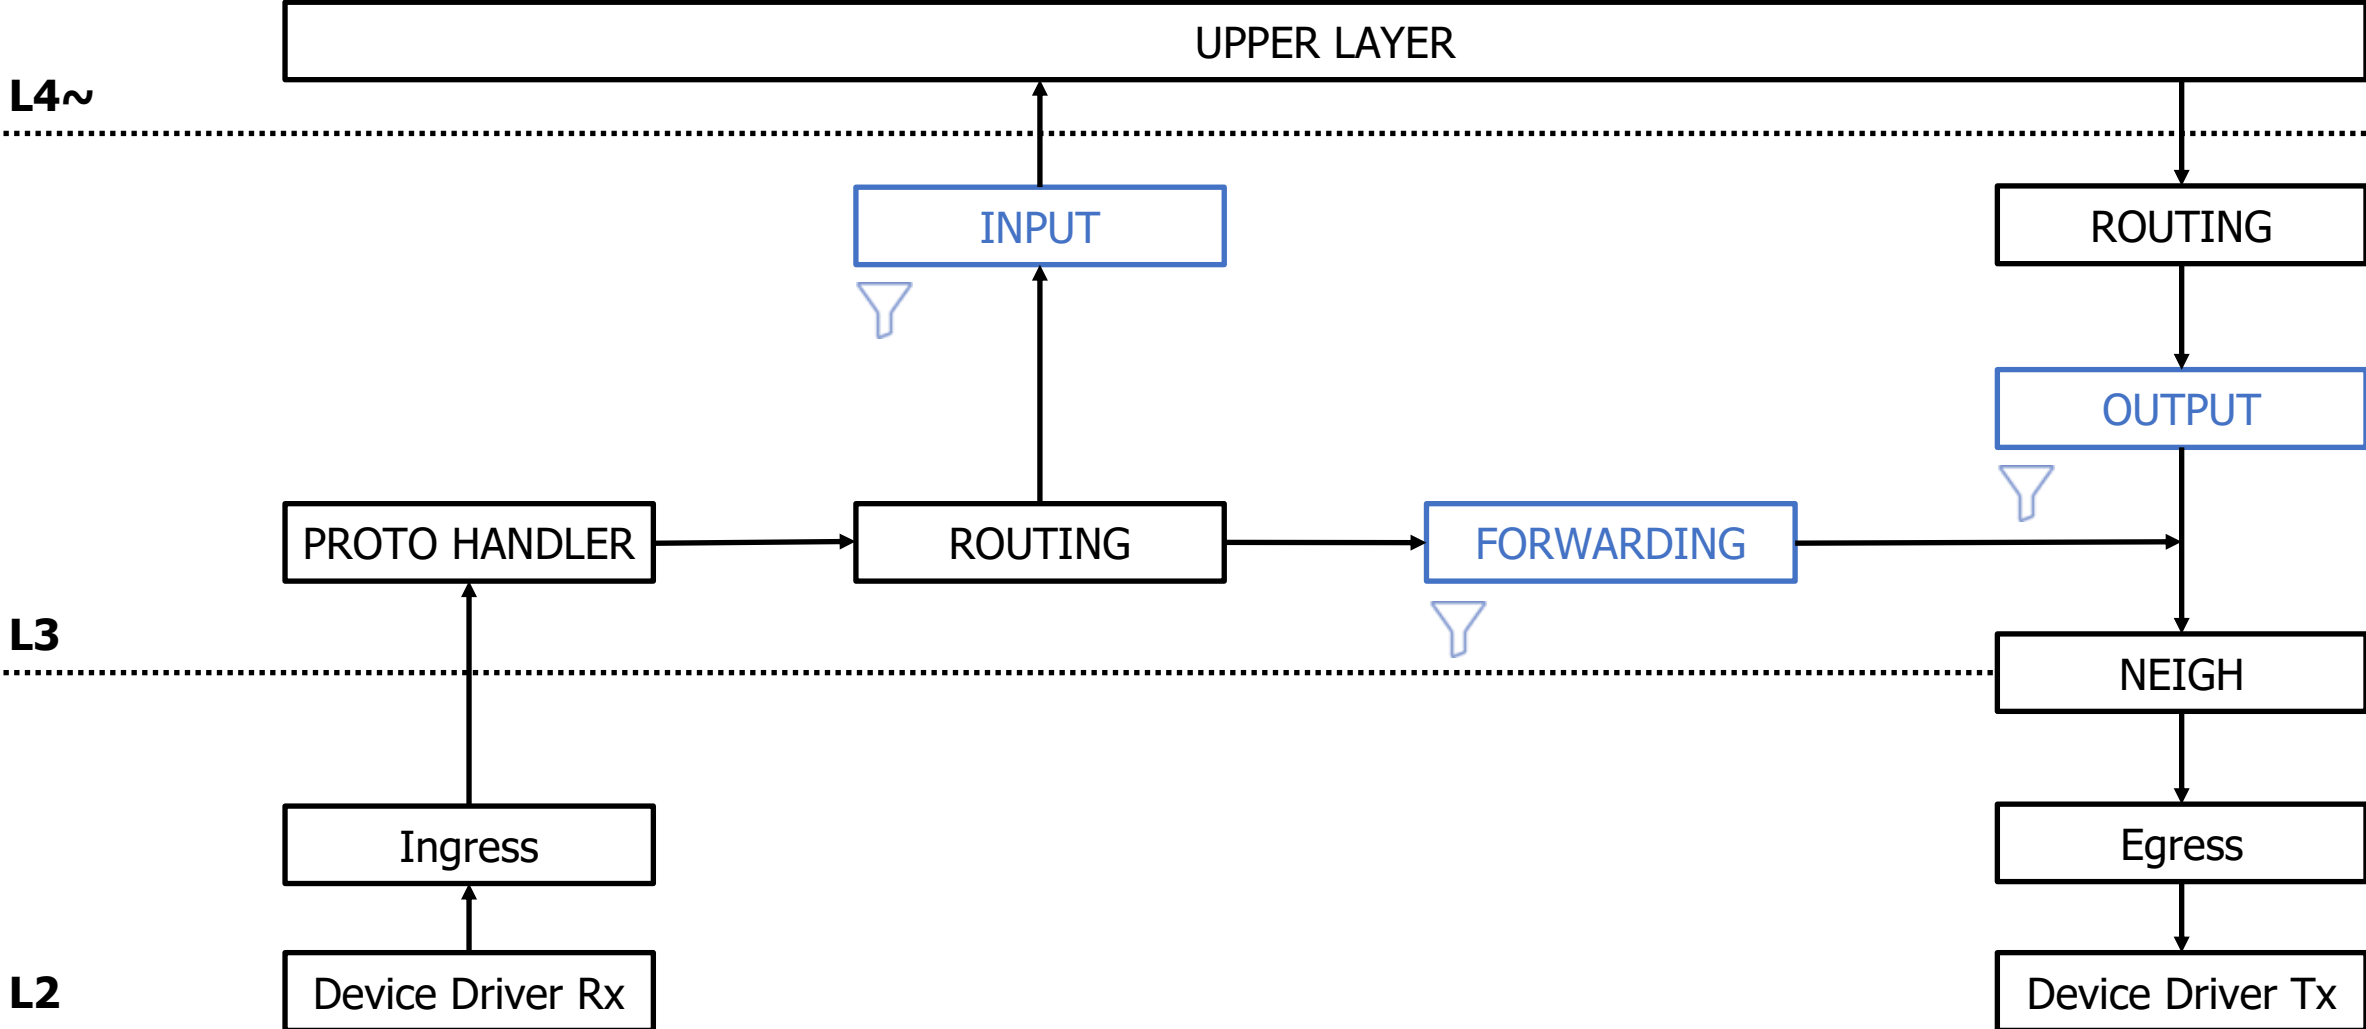
Netfilter

Linux Firewall Framework

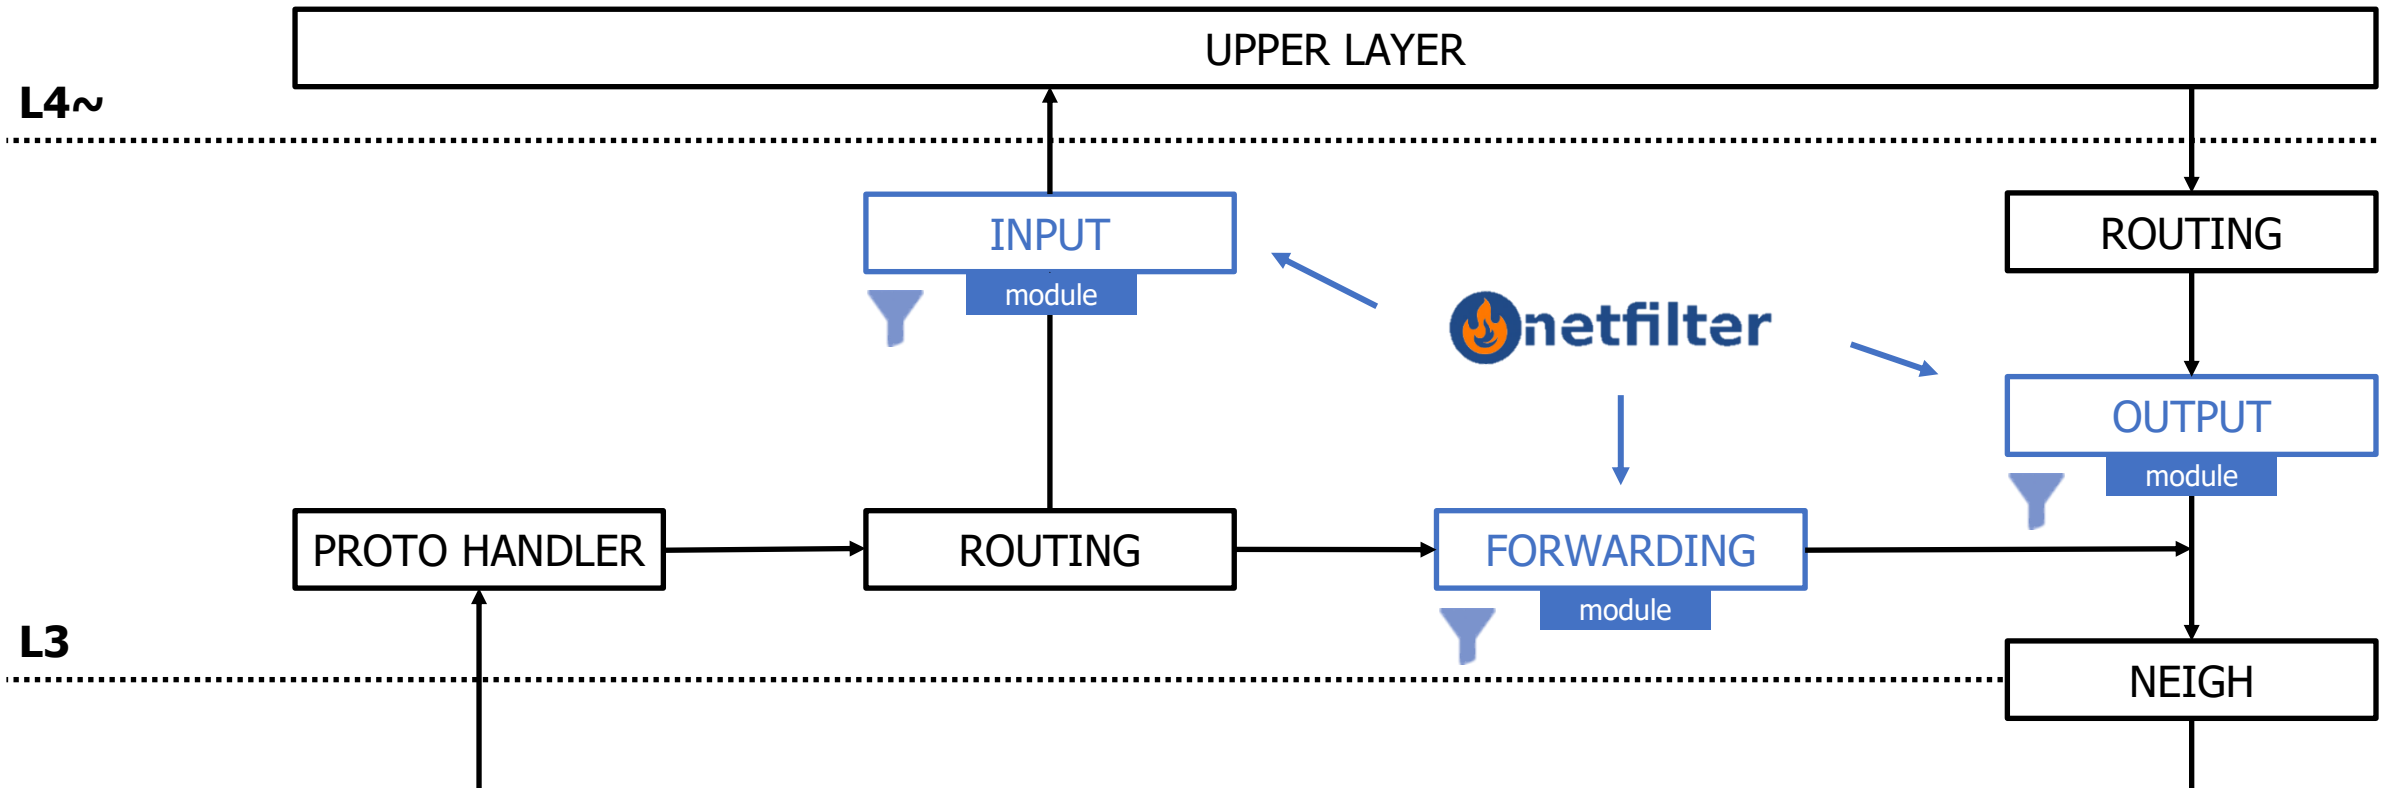
Basic Packet Rx/Tx Path



Packet Filtering Path in Kernel



Packet Filtering Path in Kernel



**Netfilter Hook에 커널 모듈을 탑재해서
패킷을 필터링할 수 있음**

ex) iptables, nftables

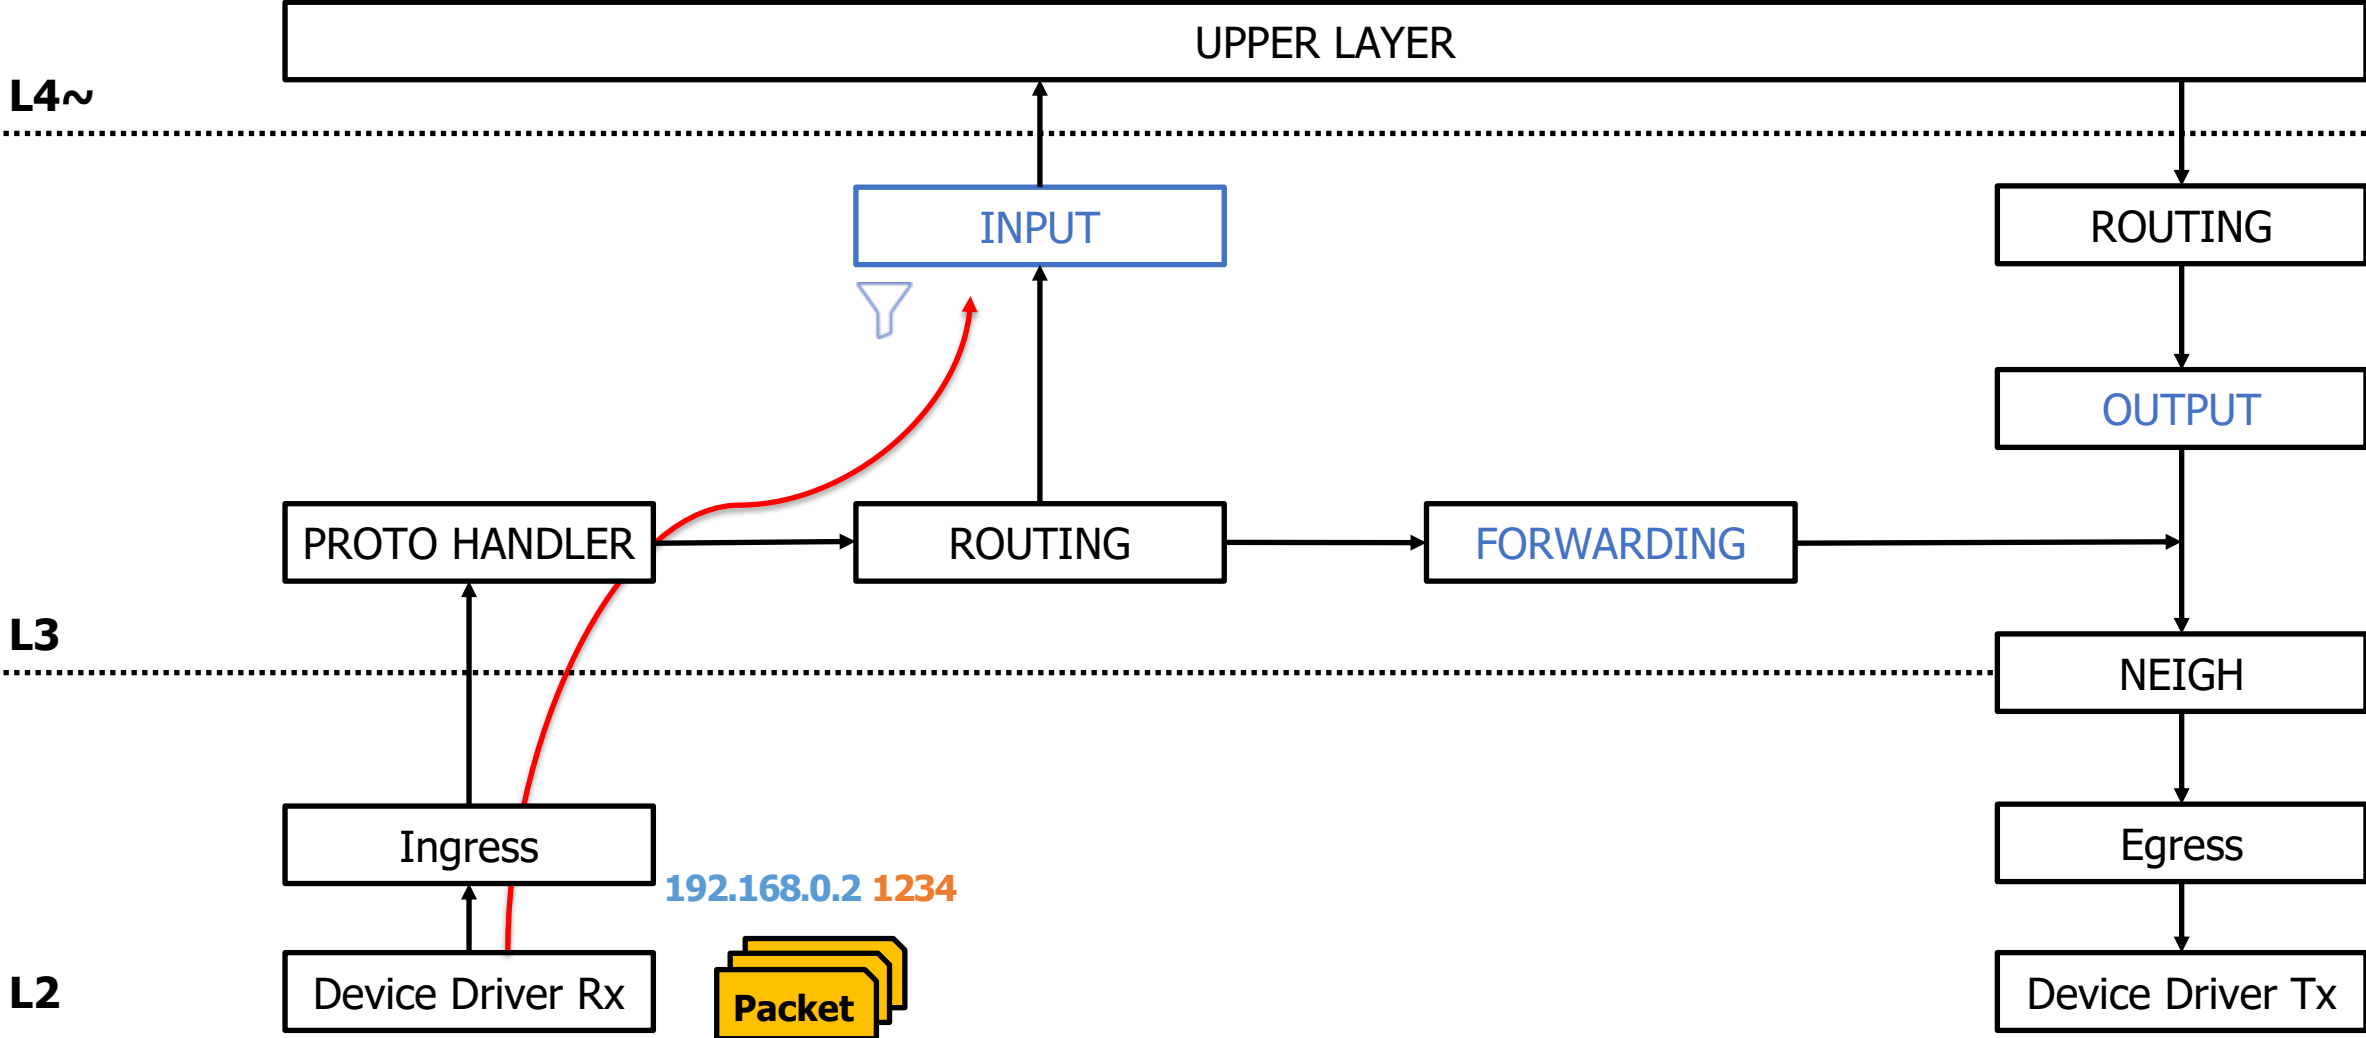
Packet Filtering Example

Netfilter (iptables) packet drop

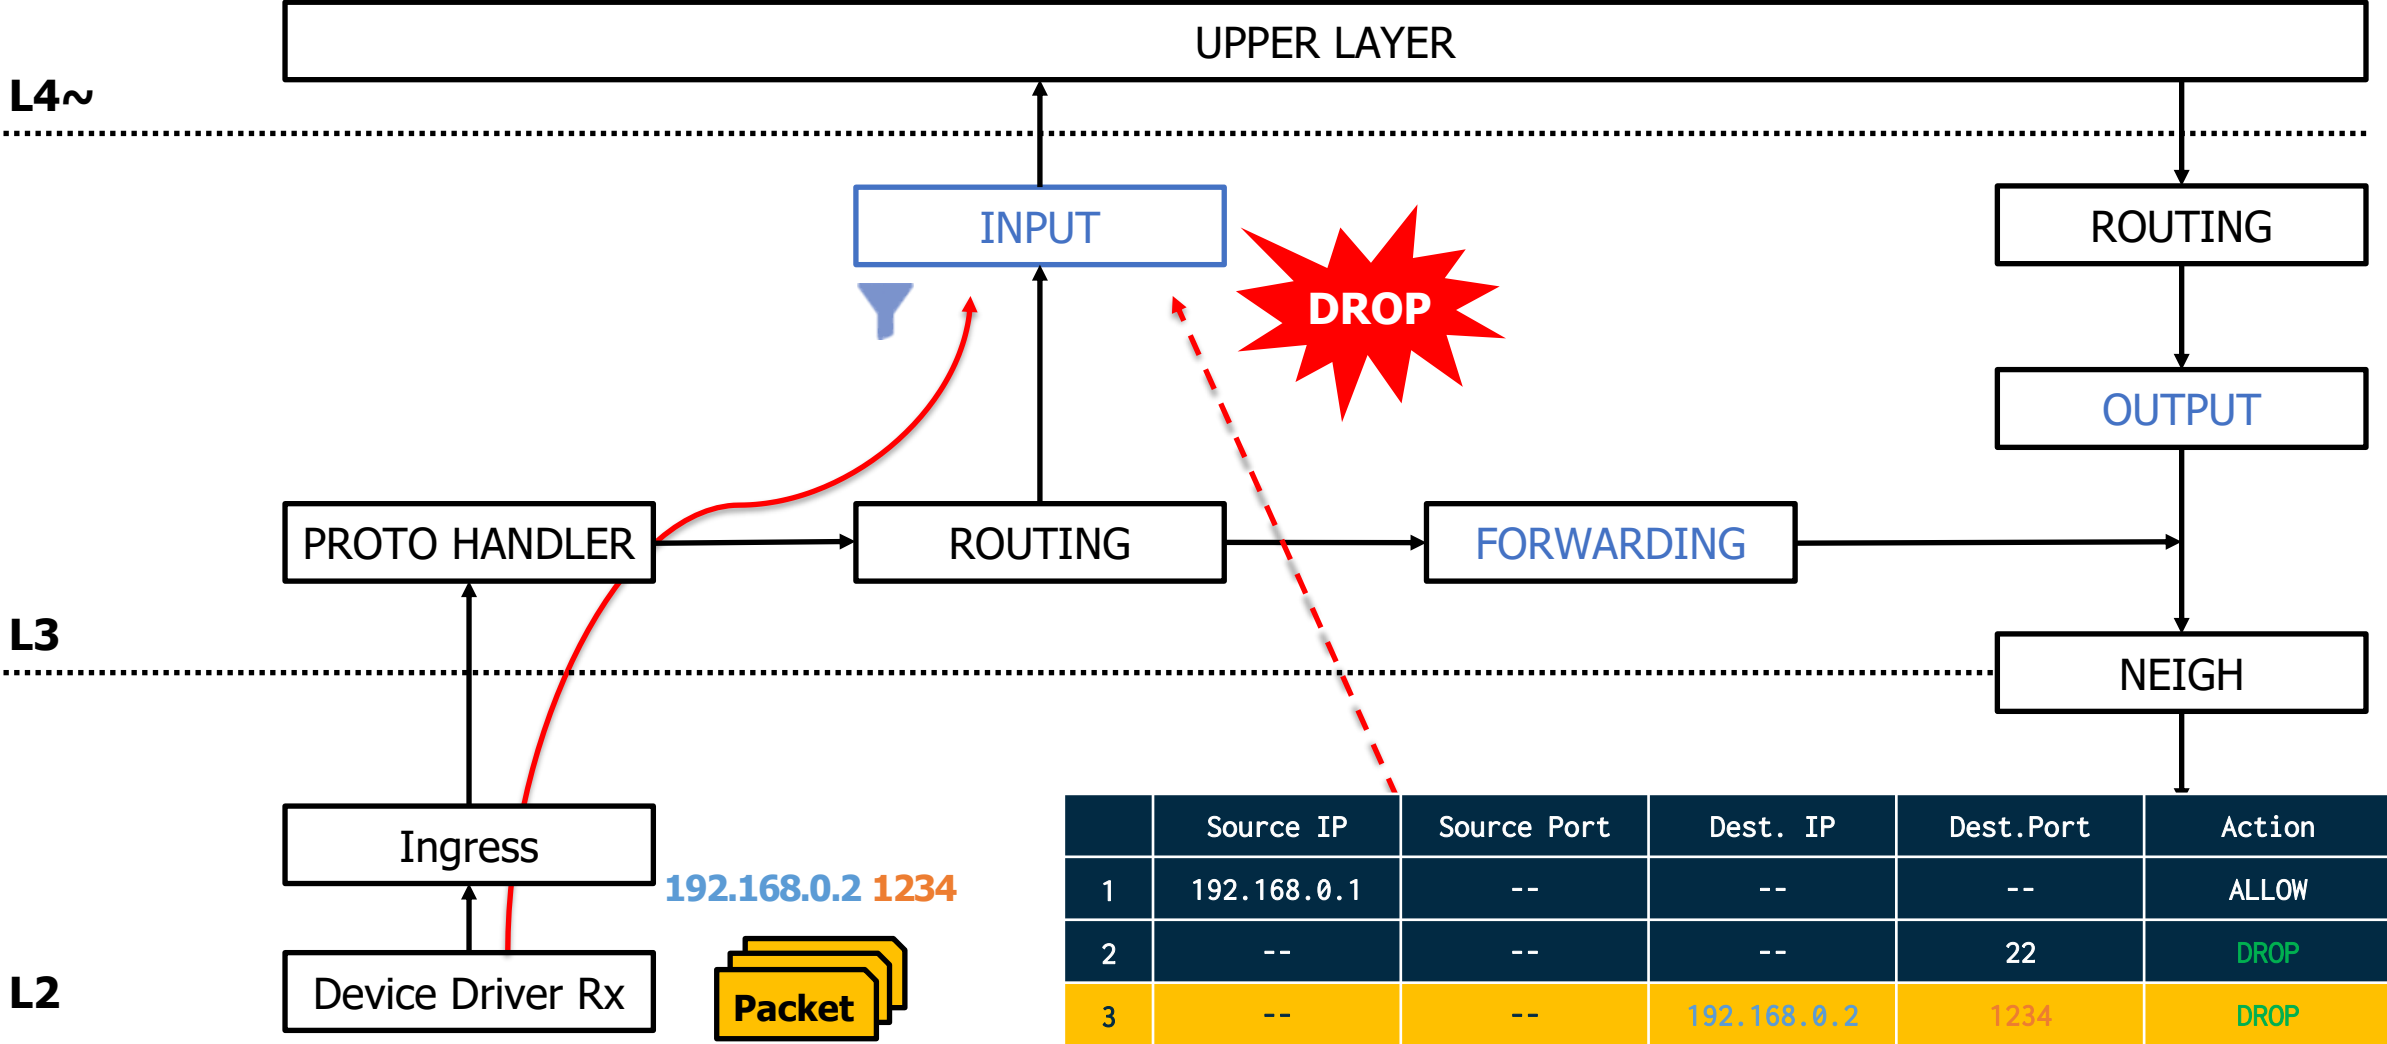
DROP UDP 192.168.0.2 1234

```
$ iptables -A INPUT -d 192.168.0.2 -p udp --dport 1234 -j DROP
```

Packet Filtering Path in Kernel



Packet Filtering Path in Kernel



NAT

Packet Sending

Routing

HSR

TCP

Packet Receiving

Packet Processing

Packet Filtering

AMT

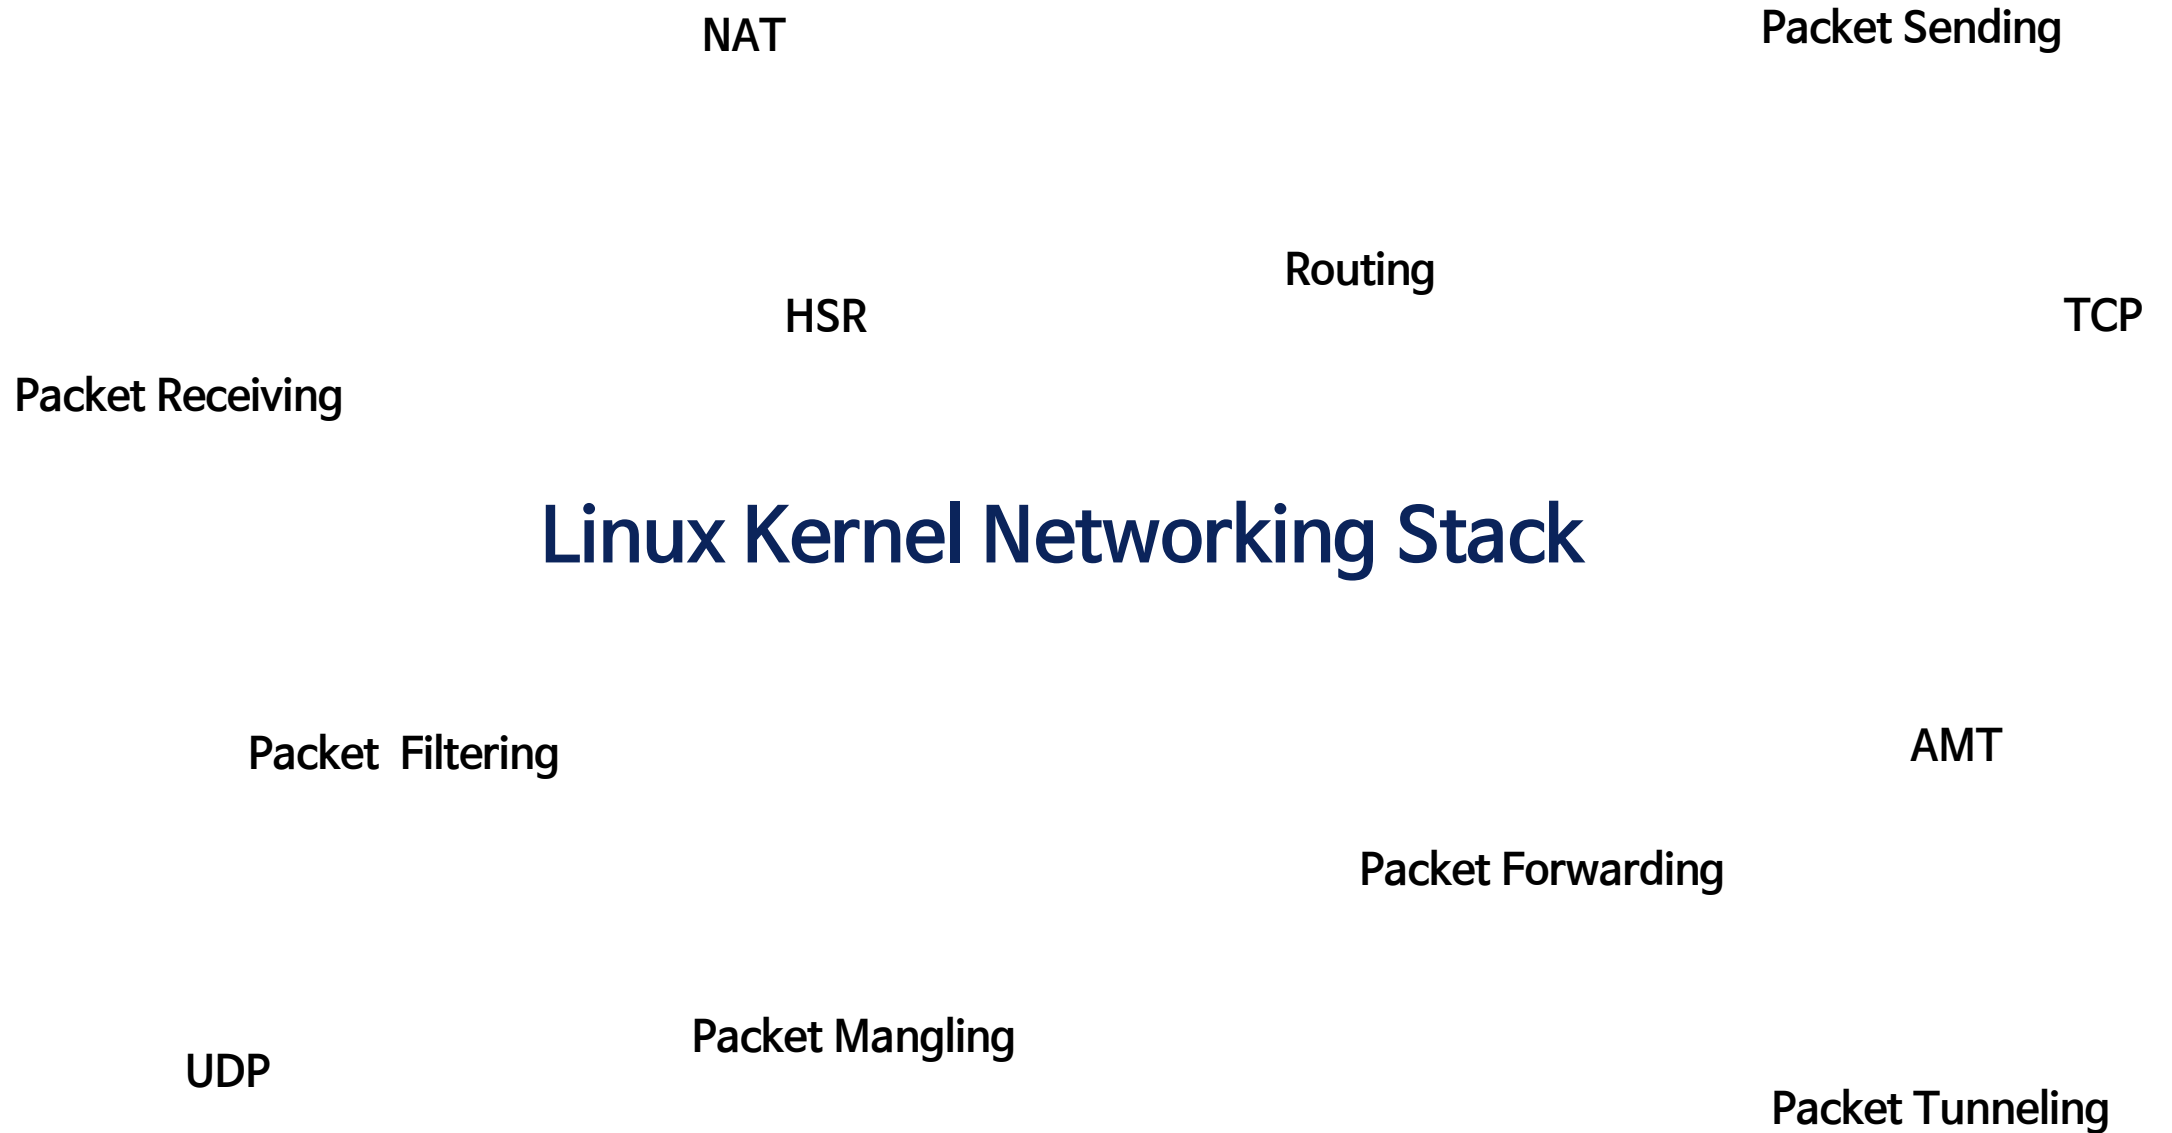
Packet Forwarding

UDP

Packet Mangling

Packet Tunneling

Linux Kernel Networking Stack



NAT

Packet Sending

Routing

HSR

TCP

Packet Receiving

Network Protocols

Packet Filtering

AMT

Packet Forwarding

UDP

Packet Mangling

Packet Tunneling

NAT

Packet Sending

Routing

TCP

HSR

Packet Receiving

Network Protocols

Packet Filtering

AMT

Packet Forwarding

UDP


Packet Mangling

Packet Tunneling

HSR

Ring Topology의 구성을 통해 중복된 패킷을 전송하여
데이터 손실 없이 통신이 되도록 하는 네트워크 프로토콜

Basic Network Packet Transfer




Critical Applications

- Fully Operational 24/7
- Require Zero-loss redundancy

**ex) Power Plant (발전소),
Transportation (교통)**

Basic Network Packet Transfer

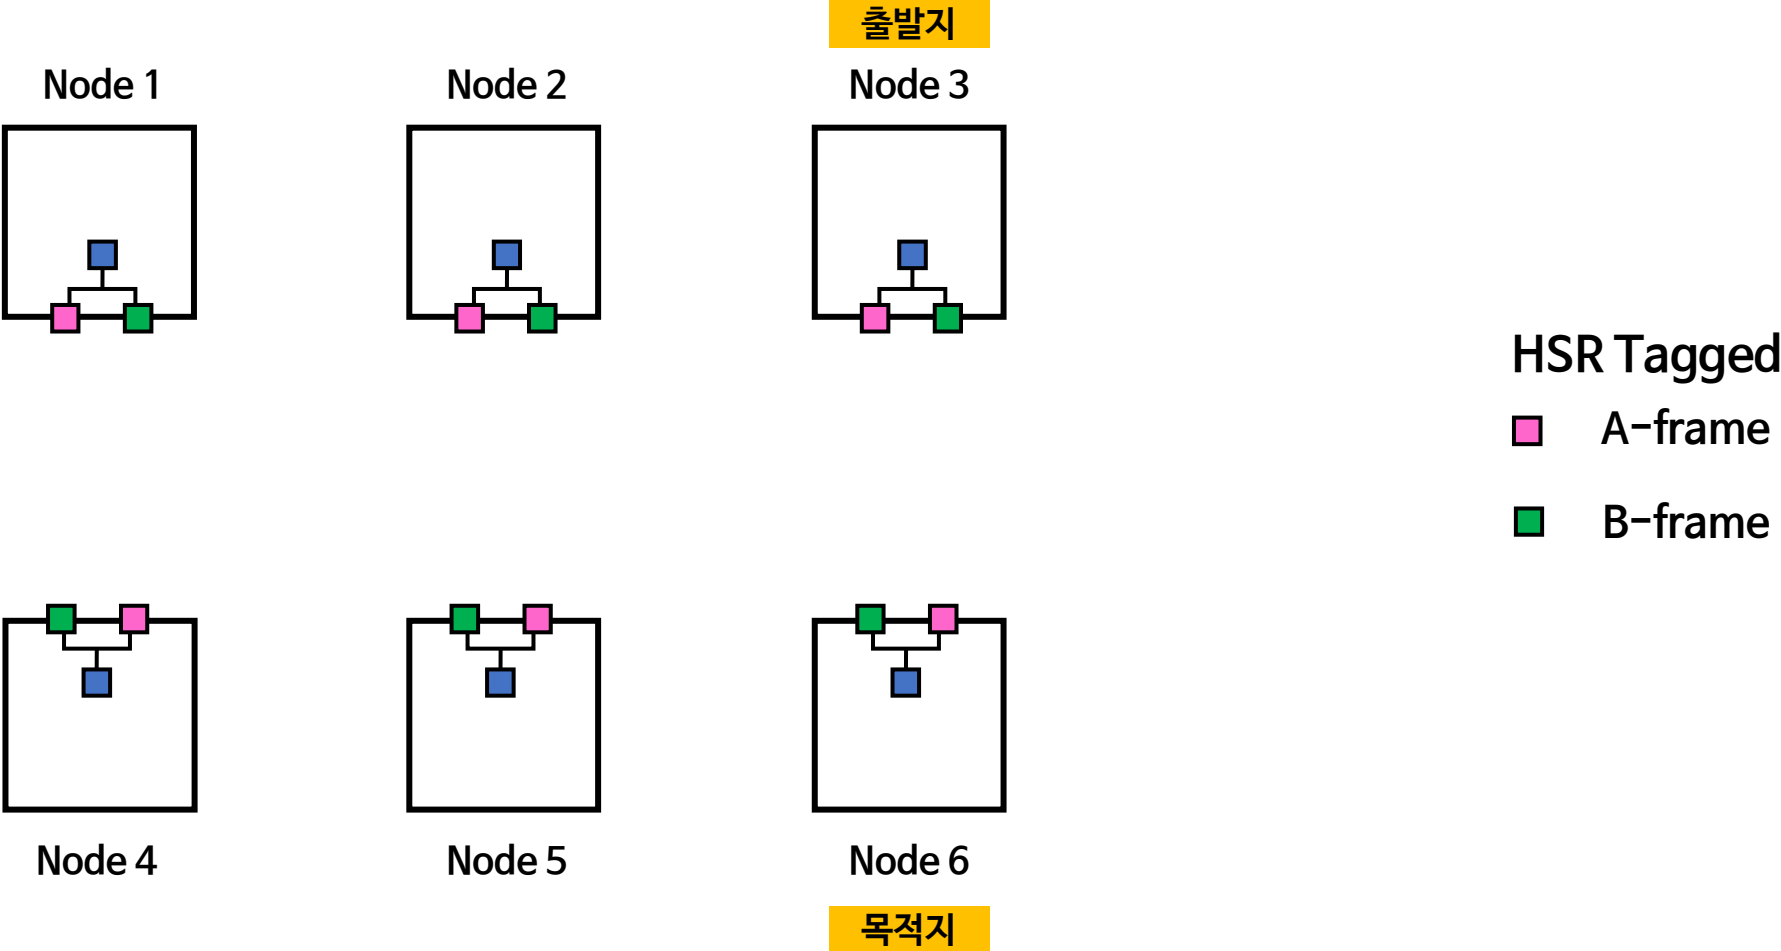


Critical Applications

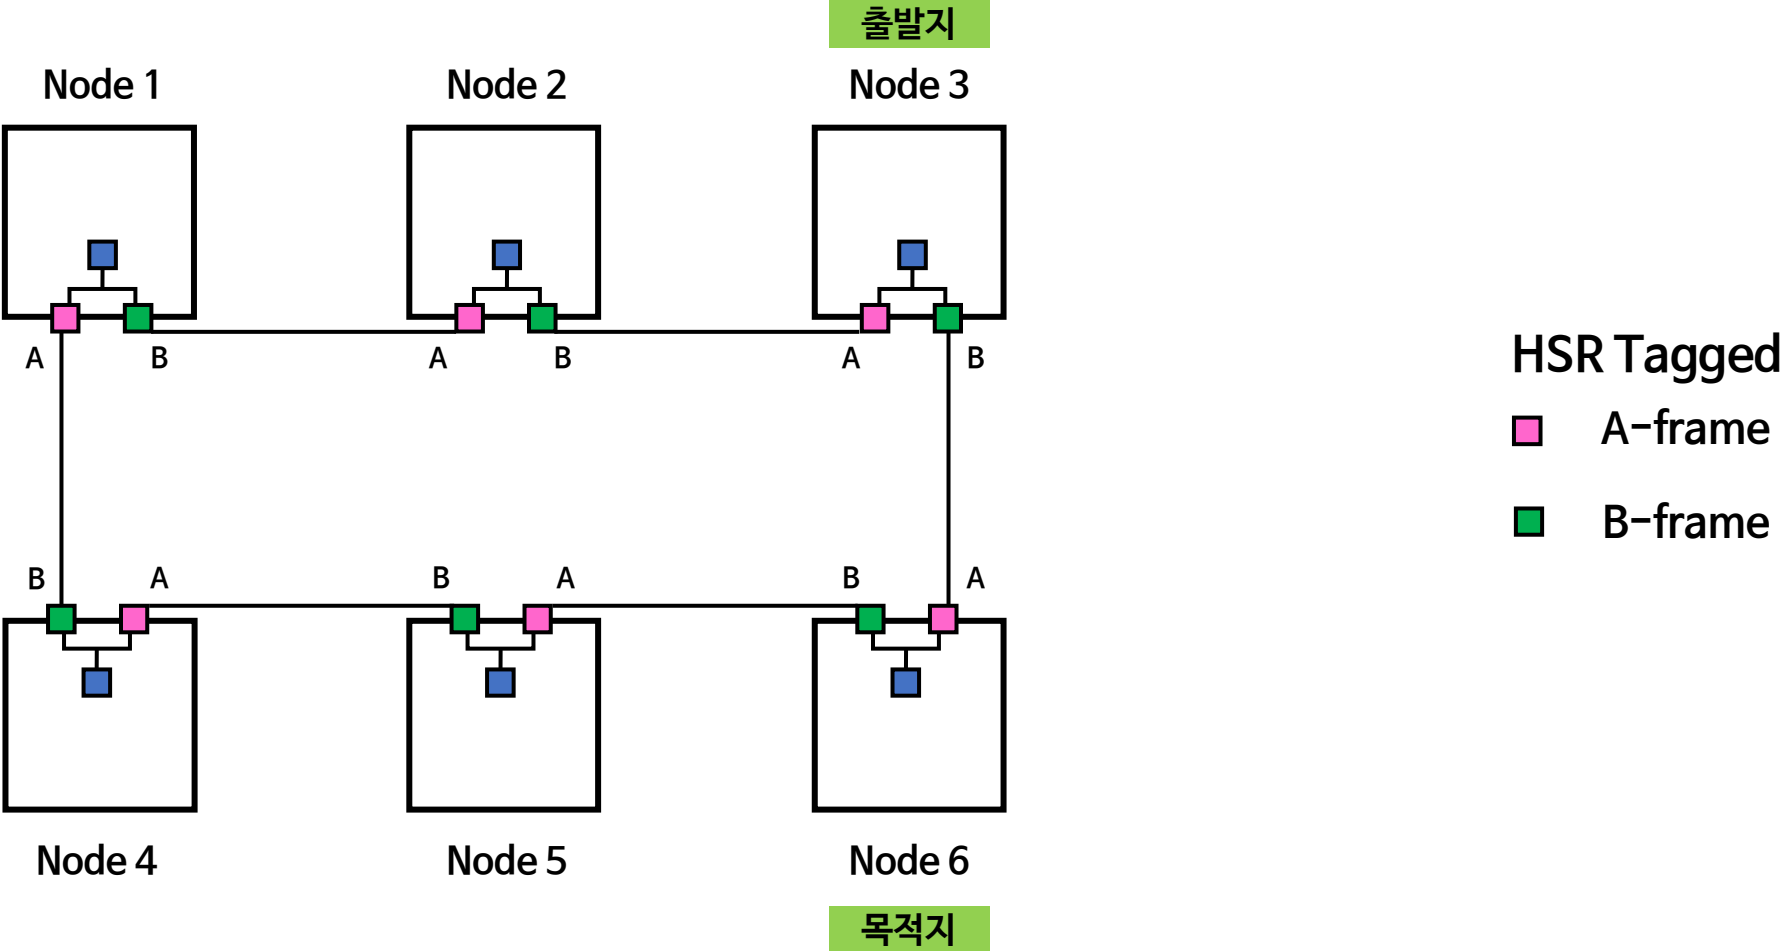
- Fully Operational 24/7
- Require Zero-loss redundancy

ex) **Power Plant (발전소),
Transportation (교통)**

High Availability Seamless Redundancy

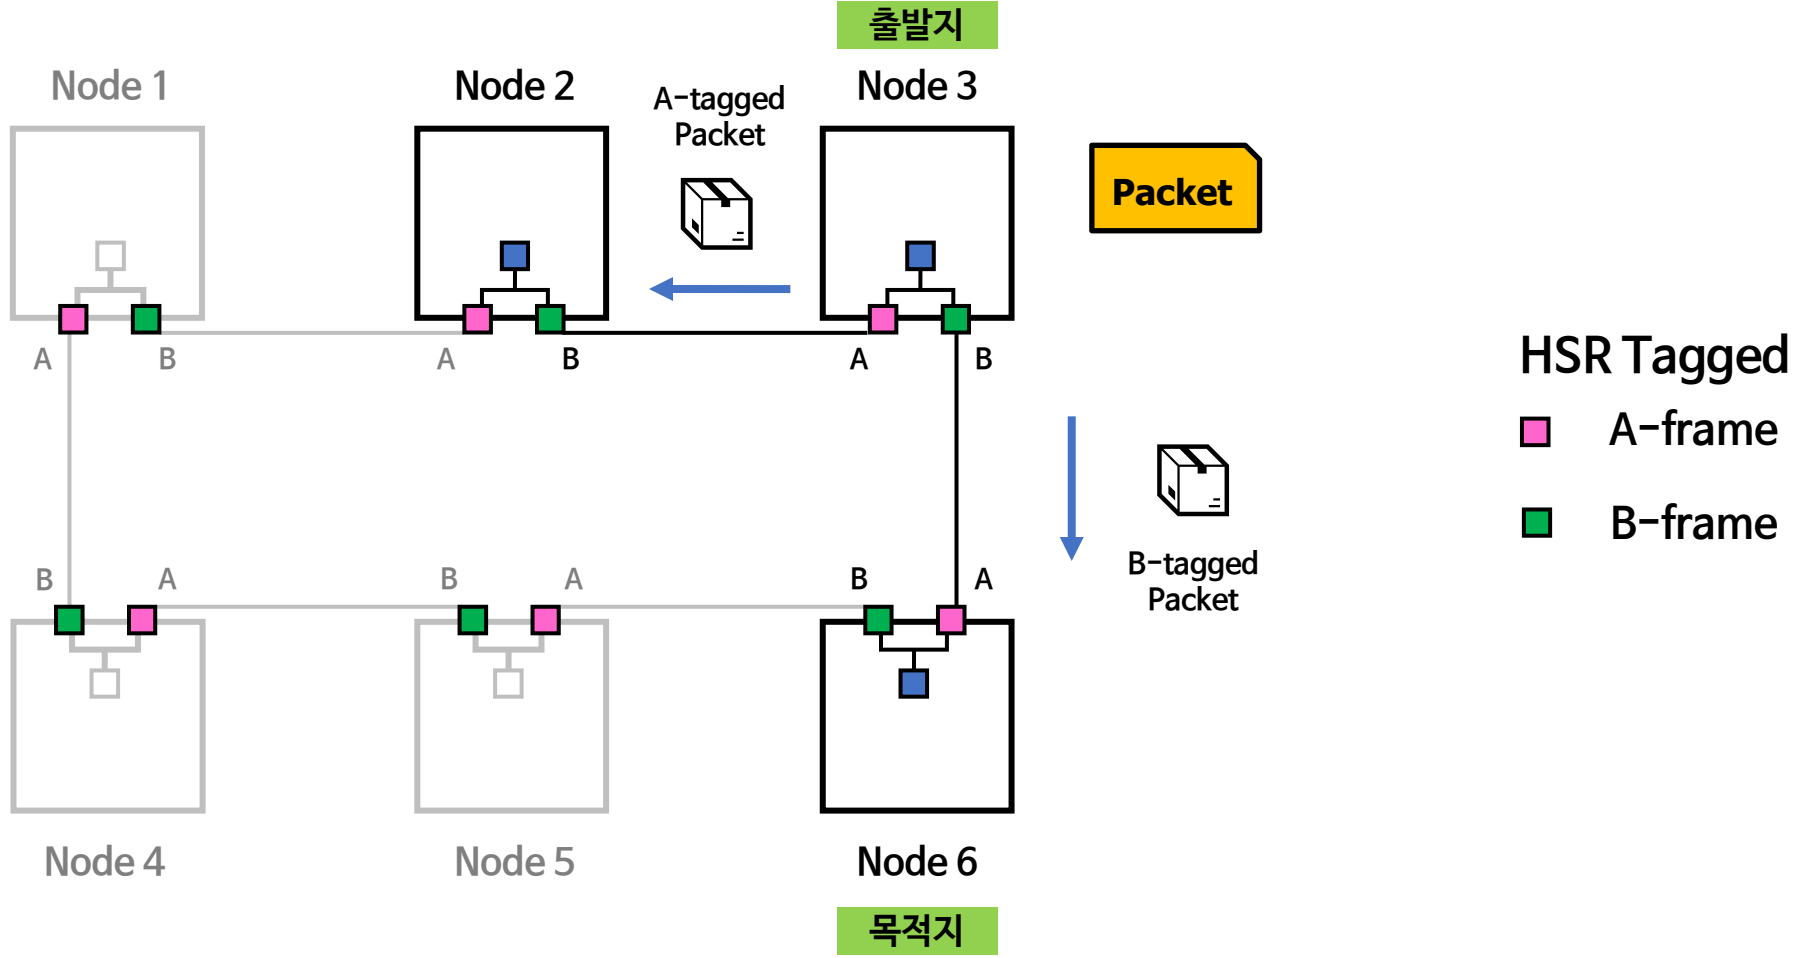


High Availability Seamless Redundancy



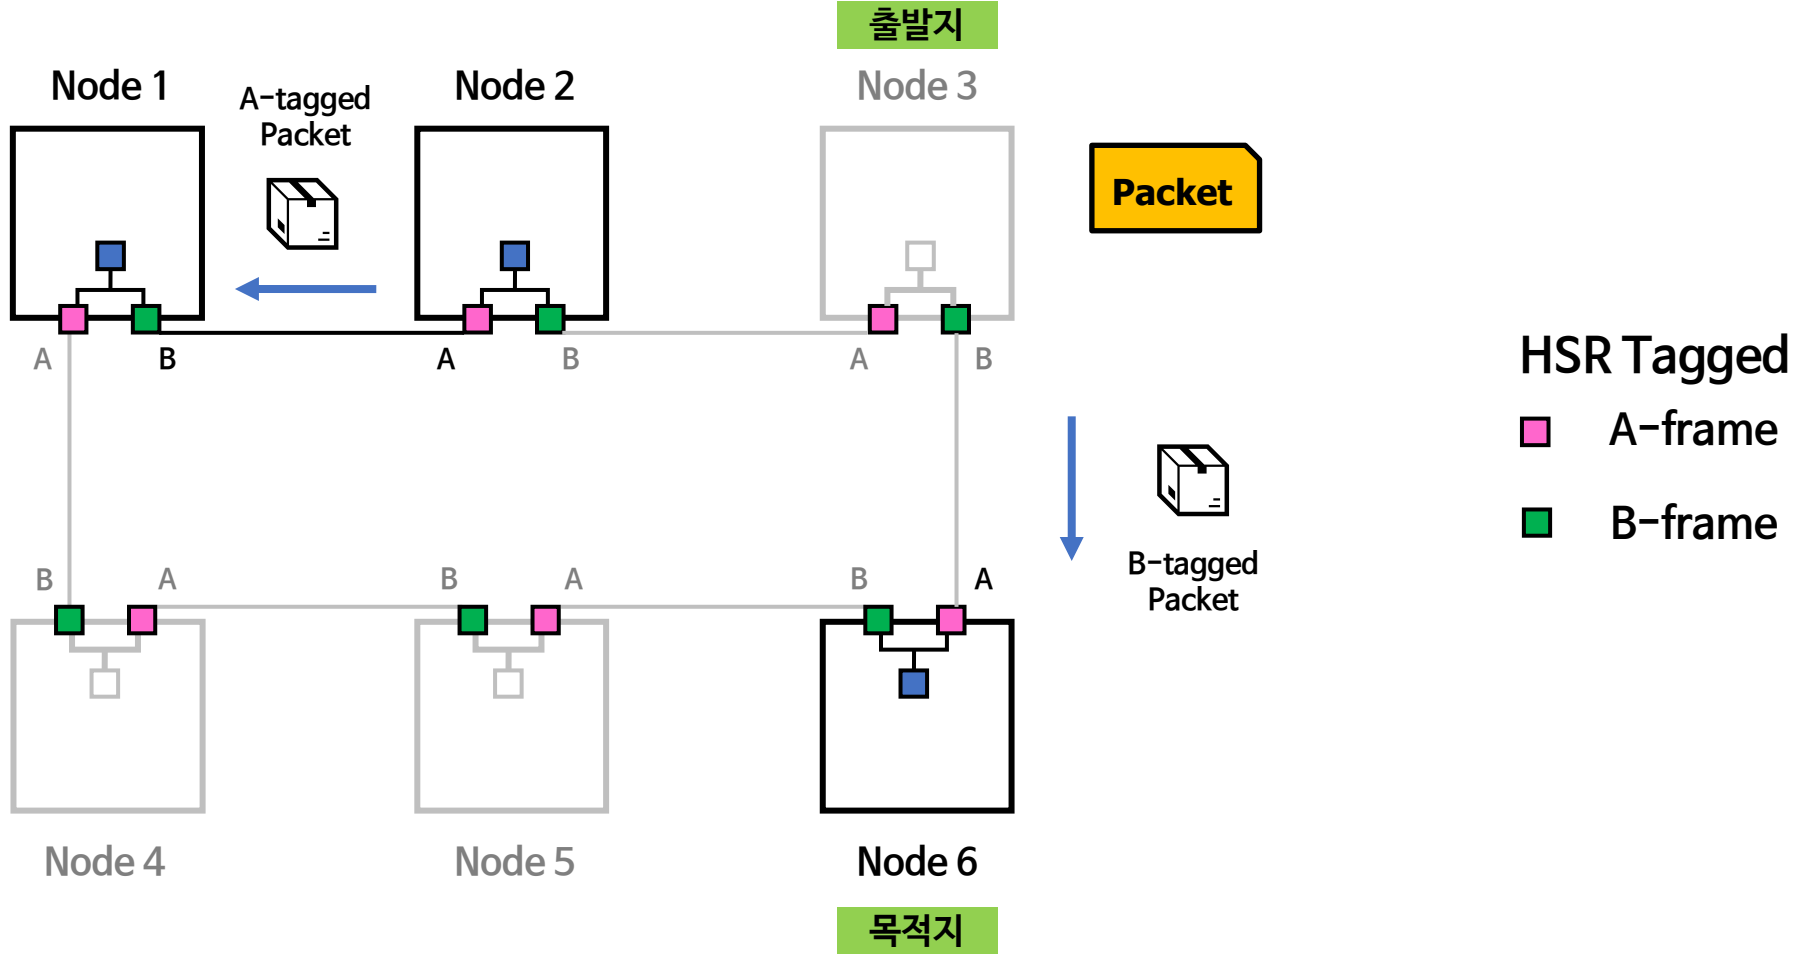
HSR Ring Topology

High Availability Seamless Redundancy



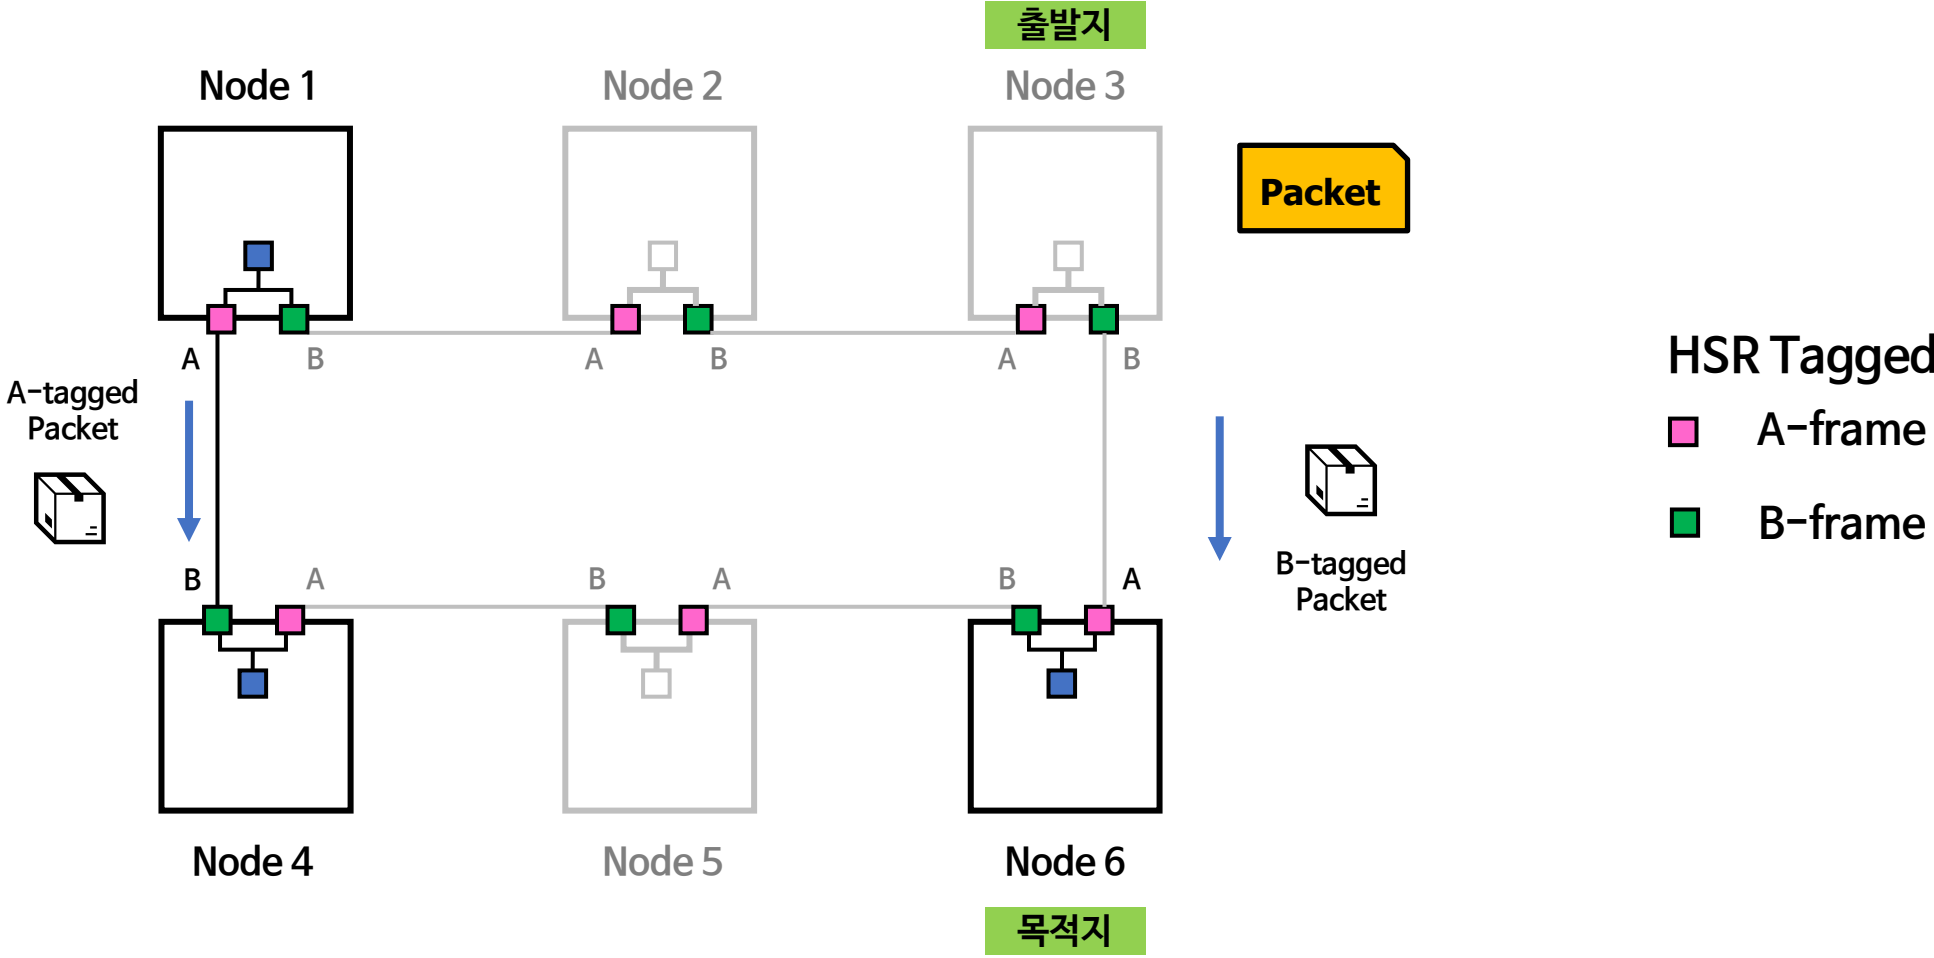
HSR Ring Topology

High Availability Seamless Redundancy



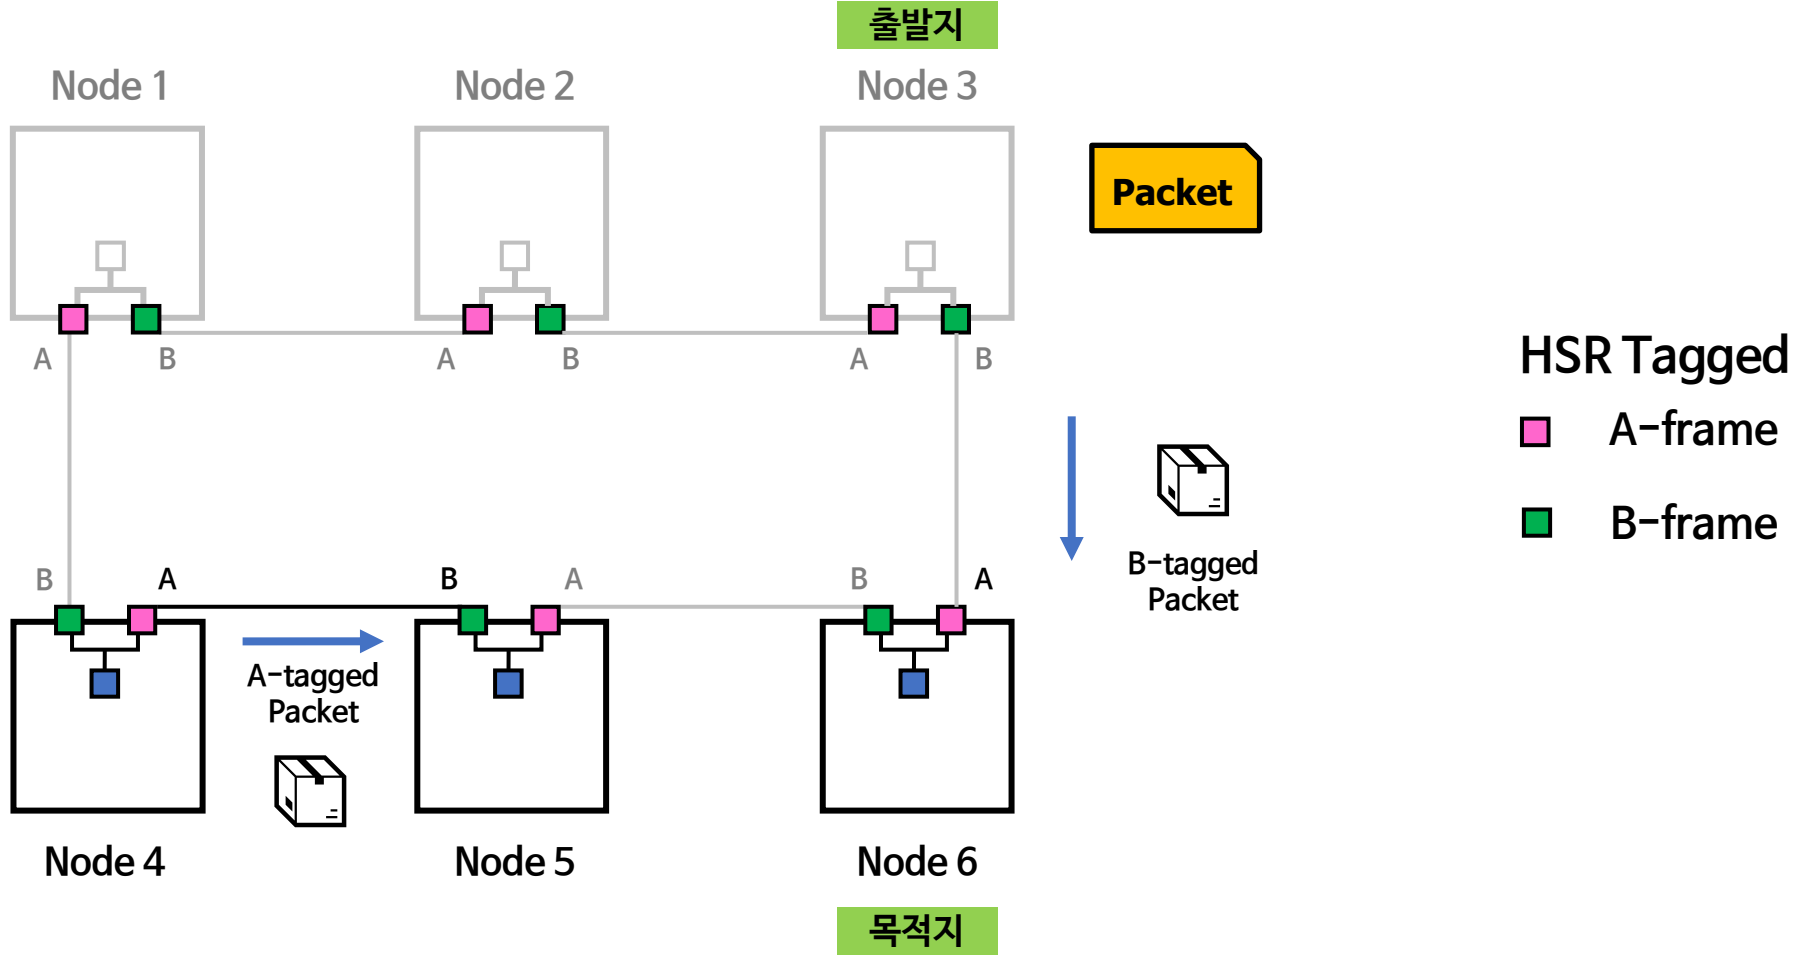
HSR Ring Topology

High Availability Seamless Redundancy



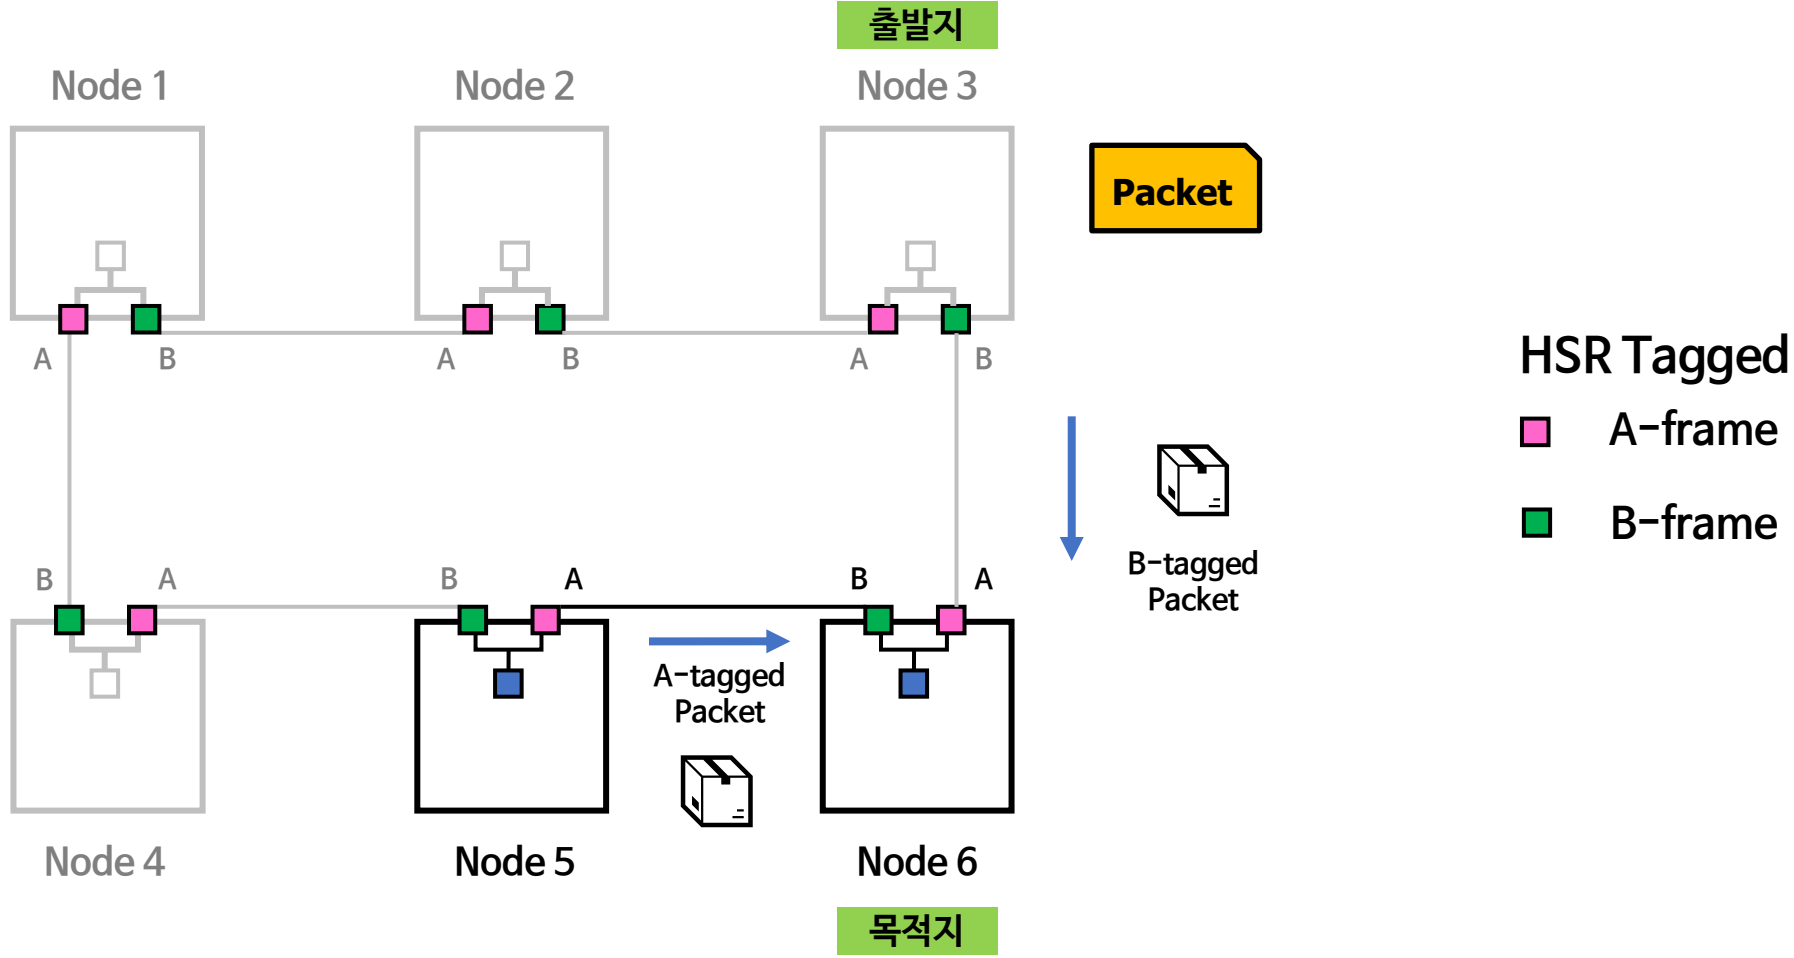
HSR Ring Topology

High Availability Seamless Redundancy



HSR Ring Topology

High Availability Seamless Redundancy



HSR Ring Topology

NAT

Packet Sending

Routing

TCP

HSR

Packet Receiving

Network Protocols

Packet Filtering

AMT

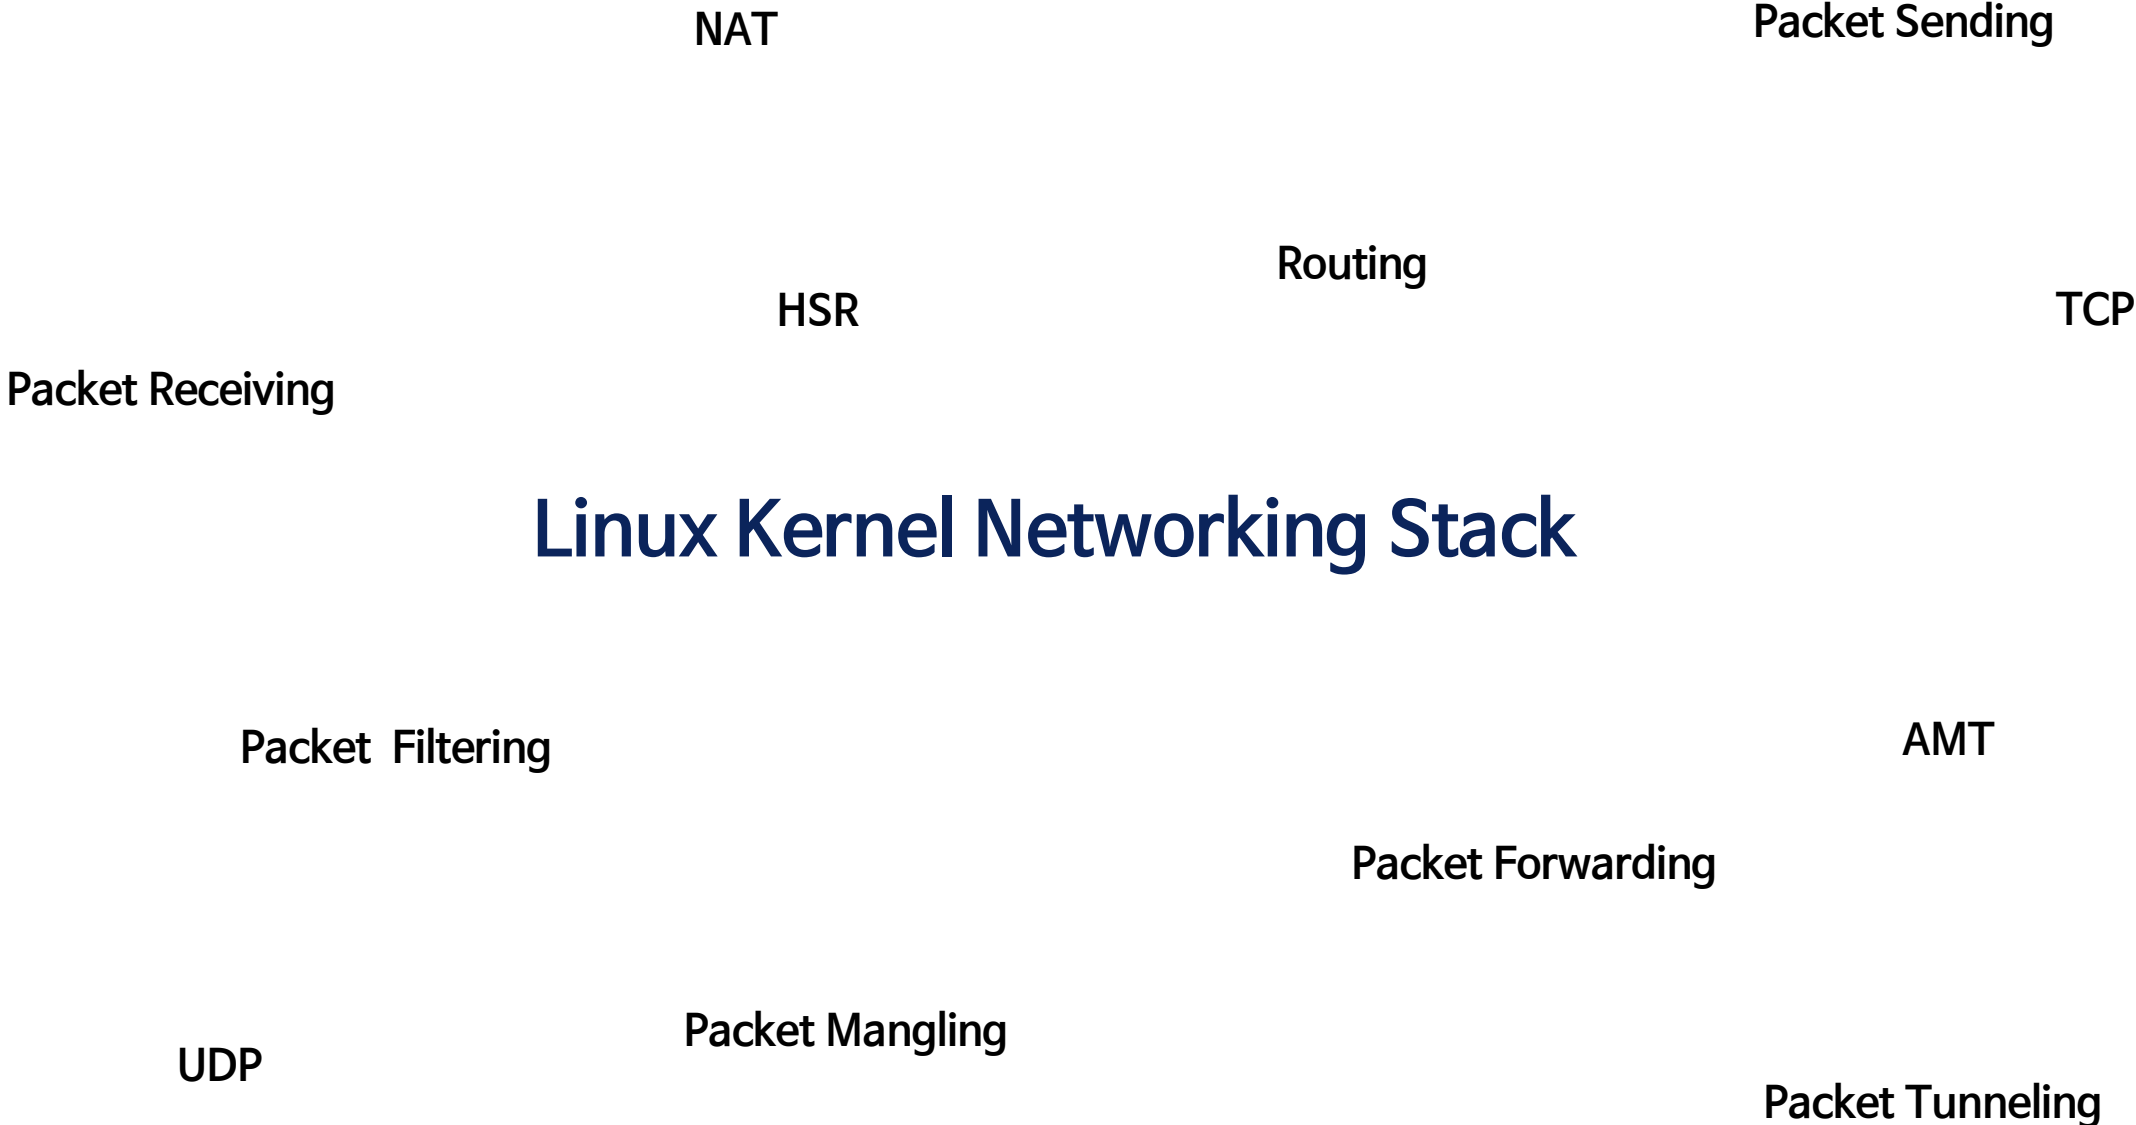
Packet Forwarding

UDP

Packet Mangling

Packet Tunneling

Linux Kernel Networking Stack




**지금부터 리눅스 커널 컨트리뷰션을 위해
리눅스 커널 생태계를 알아보시다.**

2

Kernel 생태계

Linux Kernel은 Linux 운영 체제(OS)의 핵심 구성 요소

컴퓨터 하드웨어와 프로세스를 연결하는 핵심 인터페이스

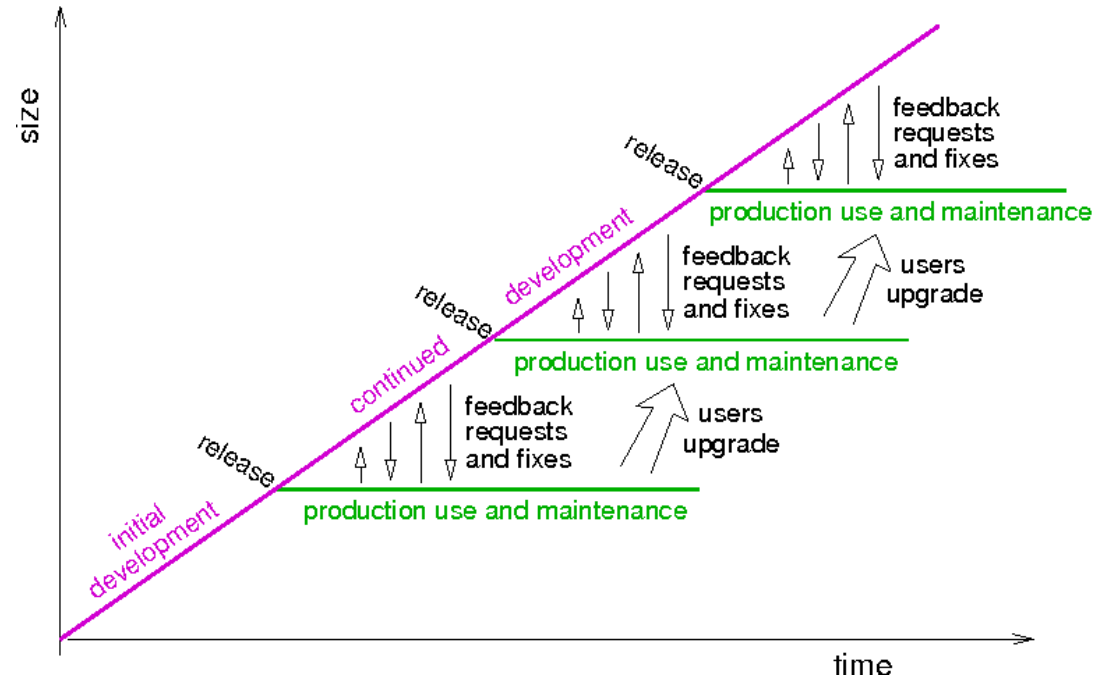


Linux Kernel 마스코트 Tux

Linux Kernel 개발

전 세계 개발자들이 참여하는
최대 규모의 오픈소스 프로젝트

개발은 **패치 (Patch)** 라는 단위로 이루어짐



<https://www.semanticscholar.org/paper/Perpetual-development%3A-A-model-of-the-Linux-kernel-Feitelson/eb7cf446e983e98c1400c8181949f038caf0c8a8/figure/1>


커널 release



10-11주(약 2개월) 마다 신규 커널 버전이 출시가 되고, stable 커널은 주 1회마다 출시

1. Mainline Kernel
2. Release Candidate (RC)
3. Stable
4. Long-term

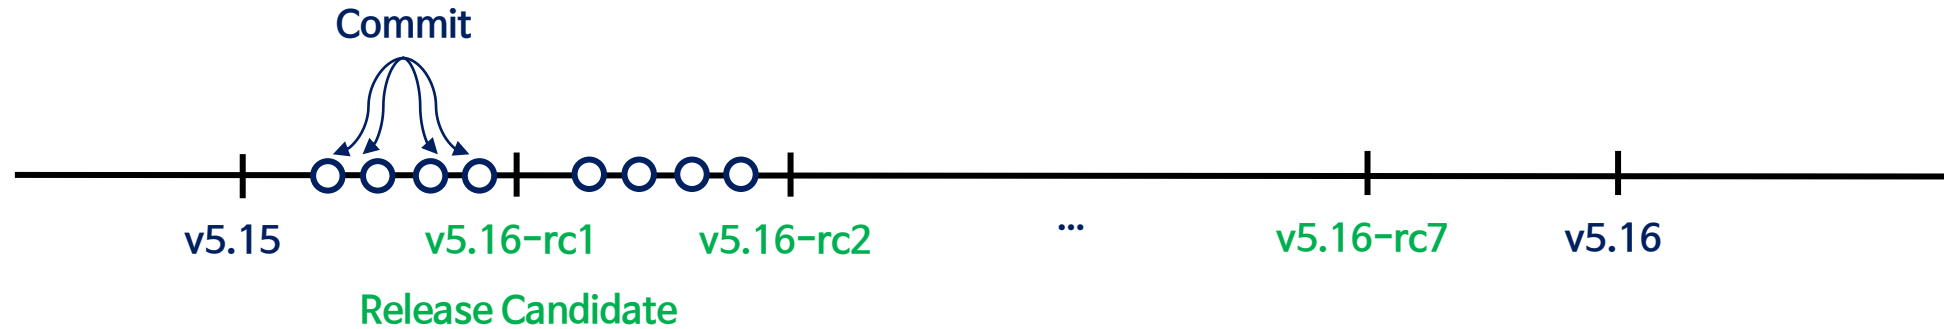
커널 release 종류



1. Mainline Kernel

- 9~10주 주기로 릴리즈되는 커널
- 새로운 기능들이 소개되고 주요 버그 fix 패치가 합쳐져 출시

커널 release 종류



2. Release Candidate (RC)

- Mainline kernel의 pre-release라고도 부름
- 신규 기능을 테스트하는 커널 버전
- 이 release는 반드시 소스코드를 컴파일해서 사용해야 함
- 주로 rc7까지 진행된 뒤 Mainline Release 진행

커널 release 종류



3. Stable

- Bug fix만 진행되는 kernel release 버전 (mainline kernel backporting)
- mainline Kernel이 release 된 이후에 stable mode로 전환


커널 release 종류



4. Long-term

- stable release 버전 중 critical한 bug fix를 제공하는 kernel release 버전

Linux Kernel Release Cycle



5.15 mainline kernel version release

Linux Kernel Release Cycle




5.15 mainline kernel version release

Linux Kernel Release Cycle



5.15 mainline kernel version release

Linux Kernel Release Cycle



5.15 mainline kernel version release ->

Linux Kernel Release Cycle




5.15 mainline kernel version release -> 5.16-rc1

Linux Kernel Release Cycle




5.15 mainline kernel version release -> 5.16-rc1 -> 5.16-rc2

Linux Kernel Release Cycle




5.15 mainline kernel version release -> 5.16-rc1 -> 5.16-rc2 -> ...

Linux Kernel Release Cycle



5.15 mainline kernel version release -> 5.16-rc1 -> 5.16-rc2 -> ... -> 5.16-rc7

Linux Kernel Release Cycle




5.15 mainline kernel version release -> 5.16-rc1 -> 5.16-rc2 -> ... -> 5.16-rc7 -> merge window

Linux Kernel Release Cycle



<http://vger.kernel.org/~davem/net-next.html>

Linux Kernel Release Cycle




5.15 mainline kernel version release -> 5.16-rc1 -> 5.16-rc2 -> ... -> 5.16-rc7 -> merge window

* Merge window

- 2주 정도 merge window 기간을 가지며, 이 merge window 기간 동안 리눅스 토발즈가 subsystem의 maintainer에게서 코드를 pull을 받아서 합치는 작업 진행


- 이 기간이 끝나면 새로운 버전 커널 release가 된 후, 2주 주기로 <mainline kernel version 명>-rc-<rc version 명>이라는 이름으로 커널 개발이 진행된다.

Linux Kernel Release Cycle




5.15 mainline kernel version release -> 5.16-rc1 -> 5.16-rc2 -> ... -> 5.16-rc7 -> merge window -> 5.16 mainline kernel version release

Linux Kernel Release Cycle



5.15 mainline kernel version release -> 5.16-rc1 -> 5.16-rc2 -> ... -> 5.16-rc7 -> merge window -> 5.16 mainline kernel version release

Linux Kernel Release Cycle



5.15 mainline kernel version release -> 5.16-rc1 -> 5.16-rc2 -> ... -> 5.16-rc7 -> merge window -> 5.16 mainline kernel version release

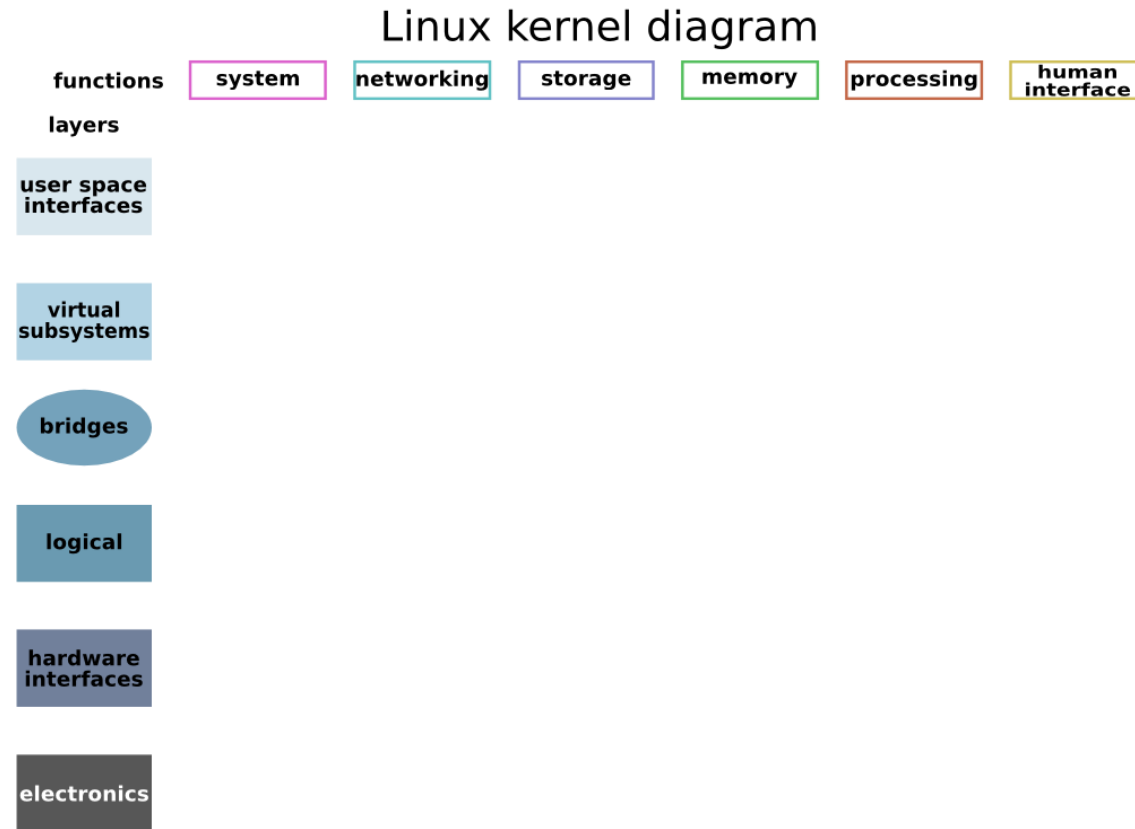
Linux Kernel Release Cycle



<http://vger.kernel.org/~davem/net-next.html>

Linux Kernel Subsystem

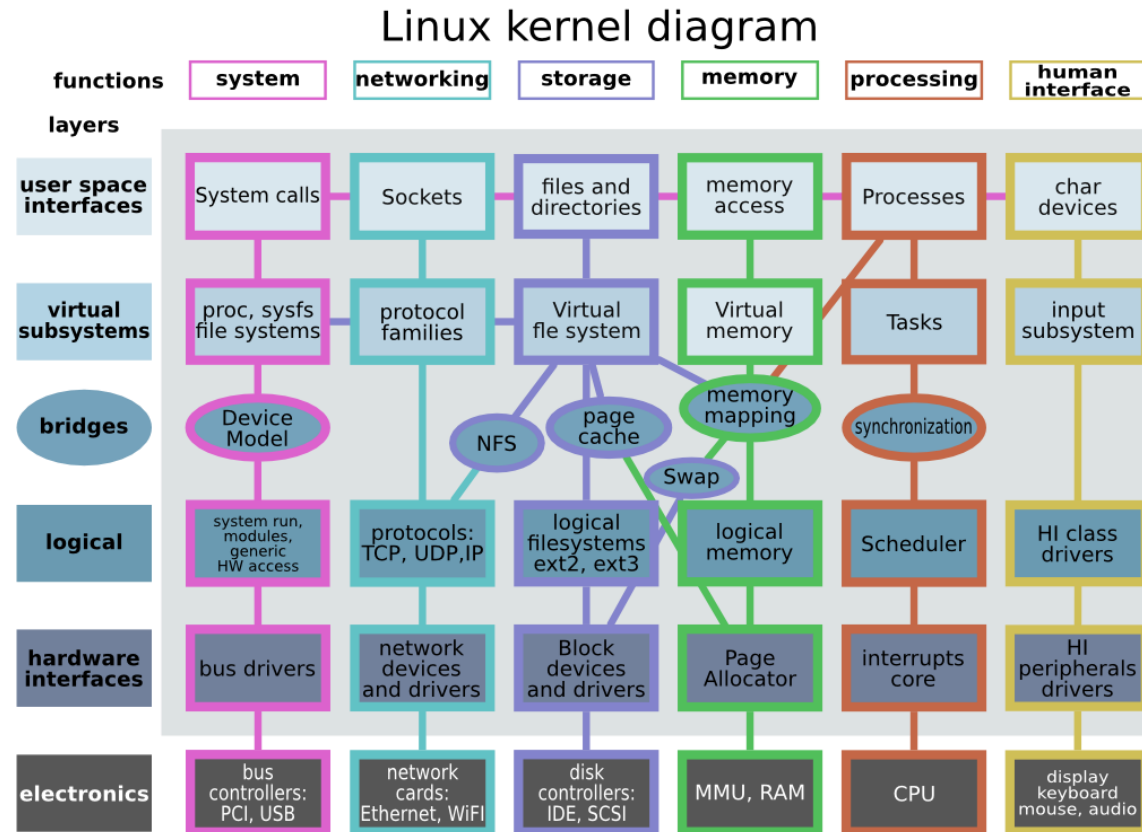
<http://www.MakeLinux.net/kernel/diagram>



커널 내부에는 **수많은 서브시스템이** 있고,

Linux Kernel Subsystem

<http://www.MakeLinux.net/kernel/diagram>



커널 내부에는 수많은 서브시스템이 있고,
서브시스템별 개발된 기능들이 합쳐져 메인라인 커널로 출시

Linux Kernel Subsystem

```
linux $ tree -d -L 1
```

```
.
├── arch
├── block
├── certs
├── crypto
├── Documentation
├── drivers
├── fs
├── include
├── init
├── ipc
├── kernel
├── lib
├── LICENSES
├── mm
├── net
├── samples
├── scripts
├── security
├── sound
├── tools
├── usr
└── virt
```

Linux Kernel v5.17 git

git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git

Linux Kernel Subsystem

```
linux $ tree -d -L 1
├── arch
├── block          => 블록 디바이스 서브시스템
├── certs
├── crypto
├── Documentation
├── drivers        => 드라이버 서브시스템
├── fs             => 파일 시스템 서브시스템
├── include
├── init
├── ipc            => IPC 서브시스템
├── kernel
├── lib
├── LICENSES
├── mm             => 메모리 관리 서브시스템
├── net            => 네트워크 서브시스템
├── samples
├── scripts
├── security
├── sound
├── tools
├── usr
└── virt
```

Linux Kernel v5.17 git

git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git

Kernel Tree



Kernel.org git repositories

Git repositories hosted at kernel.org (singapore)

<https://git.kernel.org/>

index


Name	Description	Owner	Idle	Links
<i>pub/scm/bluetooth</i>				
bluetooth-next.git	Unnamed repository; edit this file 'description' to name the repository.	Grokmirror user	19 months	summary log tree
bluez-hcidump.git	Bluetooth packet analyzer	Marcel Holtmann	5 years	summary log tree
bluez.git	Bluetooth protocol stack for Linux	Marcel Holtmann	20 hours	summary log tree
obexd.git	OBEX Server	Marcel Holtmann	5 years	summary log tree
sbc.git	Bluetooth low-complexity, subband codec (SBC) library	Marcel Holtmann	4 months	summary log tree
<i>pub/scm/boot</i>				
dracut/dracut.git	dracut - Initramfs generator using udev	Harald Hoyer	8 months	summary log tree
efilinux/efilinux.git	The efilinux UEFI boot loader	Matt Fleming	8 years	summary log tree
syslinux/syslinux.git	Unnamed repository; edit this file 'description' to name the repository.	Syslinux workgroup	6 years	summary log tree
<i>pub/scm/devel</i>				
pahole/pahole.git	Pahole and other DWARF utils	Arnaldo Carvalho de Melo	21 hours	summary log tree
sparse/chris/sparse.git	Chris Li's sparse repository.	Christopher Li	5 years	summary log tree
sparse/sparse-dev.git	Sparse's development tree	Luc Van Oostenryck	4 months	summary log tree
sparse/sparse-logs.git	Raw logs of warnings of Sparse running on the kernel with summary logs.	Luc Van Oostenryck	4 months	summary log tree
sparse/sparse.git	C semantic parser	Luc Van Oostenryck	4 months	summary log tree
<i>pub/scm/docs</i>				
docsko/korg.git	Source of korg.docs.kernel.org	Doc Group	8 months	summary log tree
kernel/kernel-docs.git	Kernel Documentation tree	Doc Group	9 years	summary log tree
kernel/ksmap.git	Kernel.org keysign map source	Kernel.org users	3 weeks	summary log tree
kernel/pgpkeys.git	Kernel developers PGP keys	Kernel.org Admin	3 days	summary log tree
kernel/website.git	Kernel.org website source	Website Editors	2 days	summary log tree

커널 코드는 몇몇 주요 subsystem git repository로 관리되는데
이걸 **Kernel Tree**라고 함

Kernel Tree

Mainline Kernel tree

- [Linus Torvalds](#)가 관리
- Mainline kernel tree와 RC release 소스 코드가 관리



* <https://www.devcurmudgeon.com/images/mainline-lts-ltsi-genivi-20131025.pdf#page=18>


Kernel Tree

Mainline Kernel tree

- Linus Torvalds가 관리
- Mainline kernel tree와 RC release 소스 코드가 관리

Stable Kernel tree

- [Greg Kroah-Hartman](#)가 관리
- stable release branches 소스 코드가 관리



* <https://www.devcurmudgeon.com/images/mainline-lts-ltsi-genivi-20131025.pdf#page=18>

Kernel Tree

Mainline Kernel tree

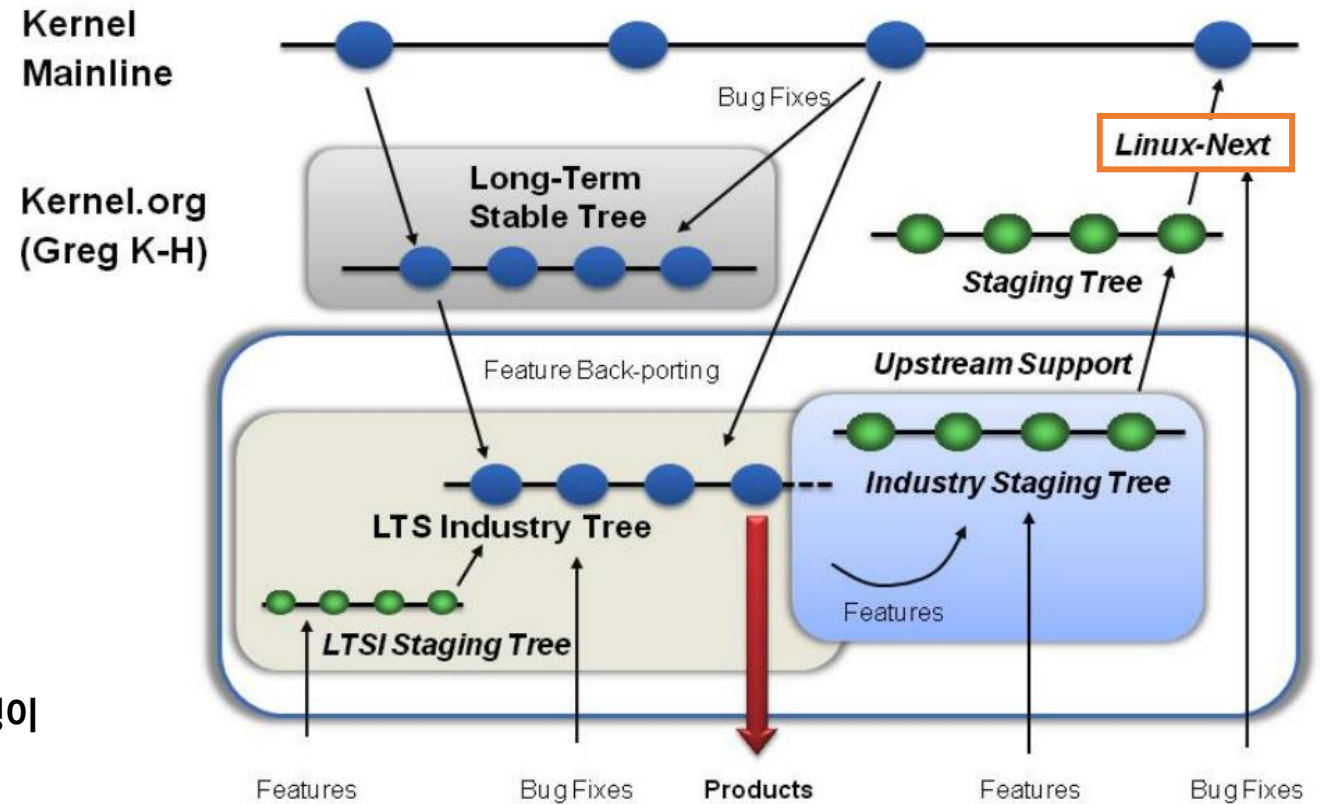
- Linus Torvalds가 관리
- Mainline kernel tree와 RC release 소스 코드가 관리

Stable Kernel tree

- Greg Kroah-Hartman가 관리
- stable release branches 소스 코드가 관리


linux-next tree

- [Stephen Rothwell](#)가 관리.
- Subsystem의 변경 사항을 주기적으로 반영하고, 해당 코드의 통합 테스트가 진행되는 곳
- 테스트가 진행된 이후 mainline kernel tree로 pull request가 진행



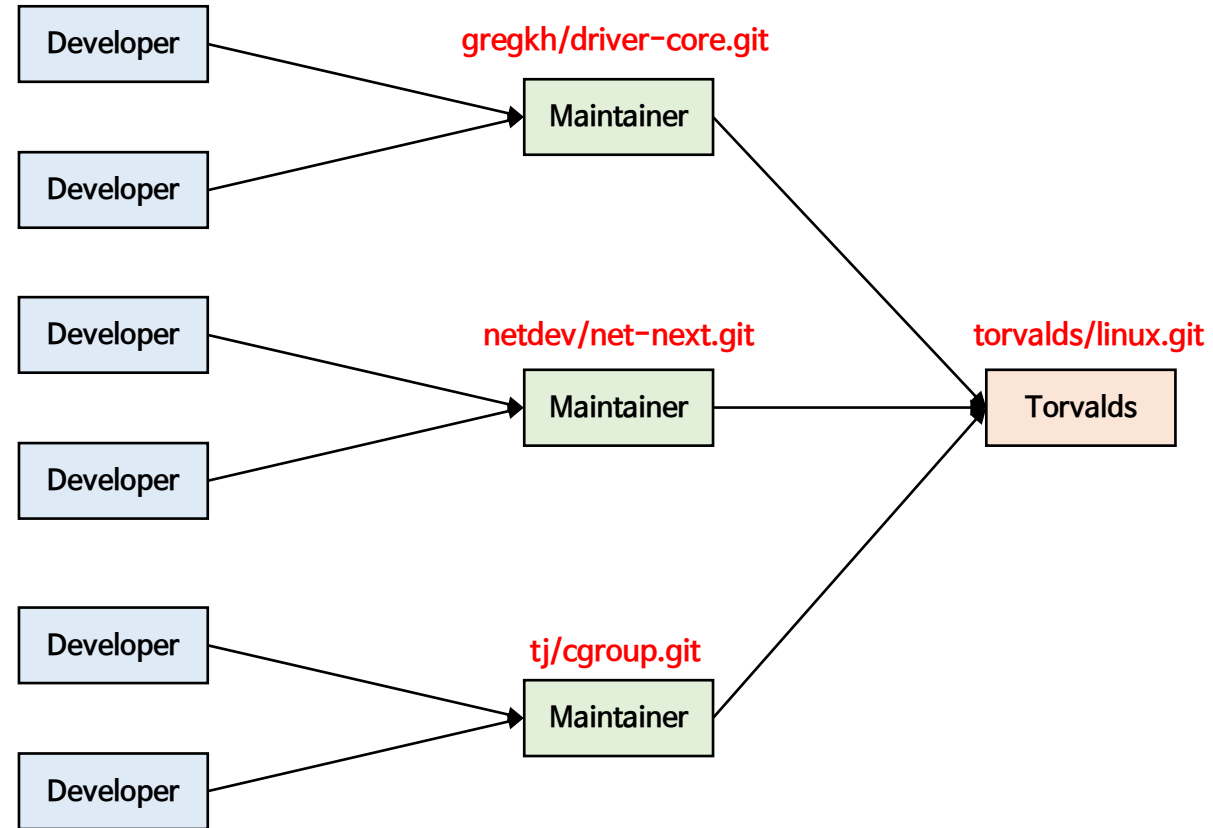
* <https://www.devcurmudgeon.com/images/mainline-lts-ltsi-genivi-20131025.pdf#page=18>

Linux Kernel Tree Hierarchy



주요 서브시스템은 독립된 tree를 가짐

Linux Kernel Tree Hierarchy



서브시스템 별 해당 tree를 관리하는 maintainer가 존재

Subsystem Maintainers



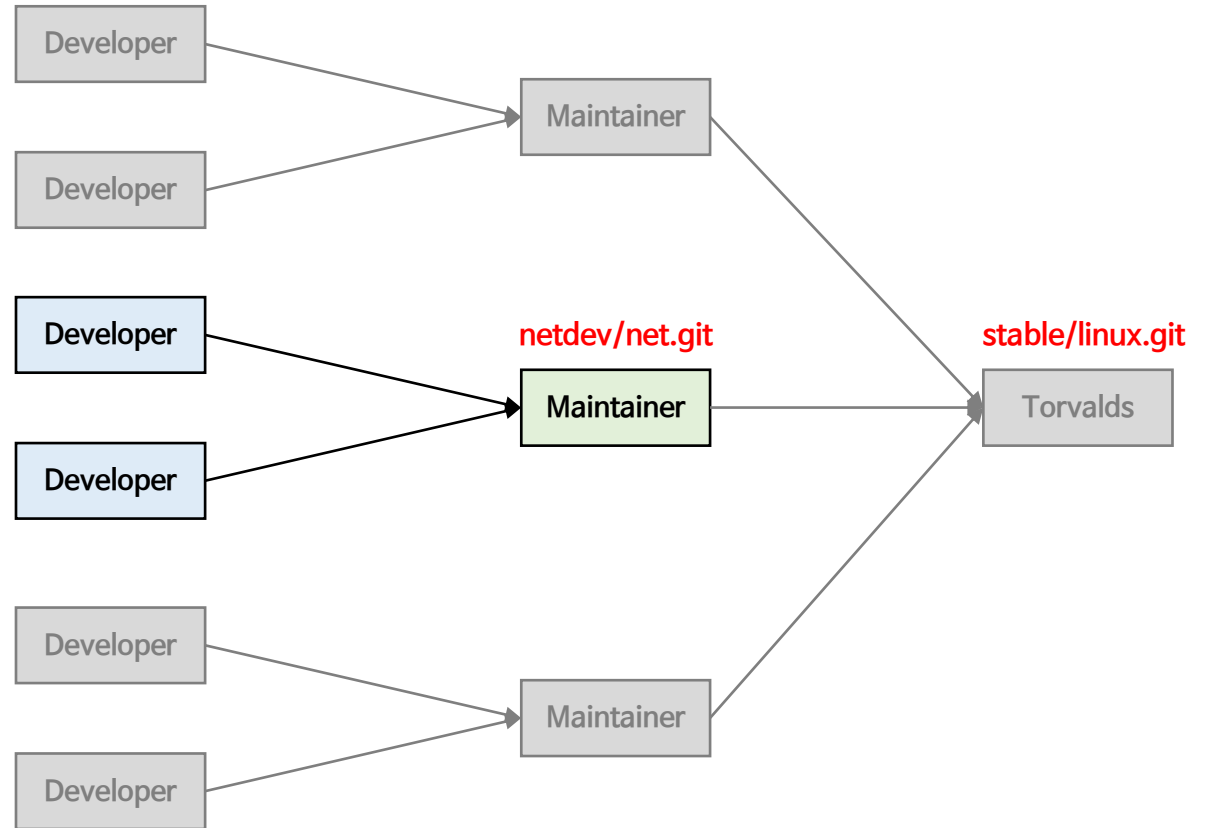
index : kernel/git/netdev/net.git

git.kernel.org/pub/scm/linux/kernel/git/netdev/net.git

Netdev Group

netdev/net.git : stable 버전 tree

기존 release 버전의 버그 fix patch만 허용



Subsystem Maintainers

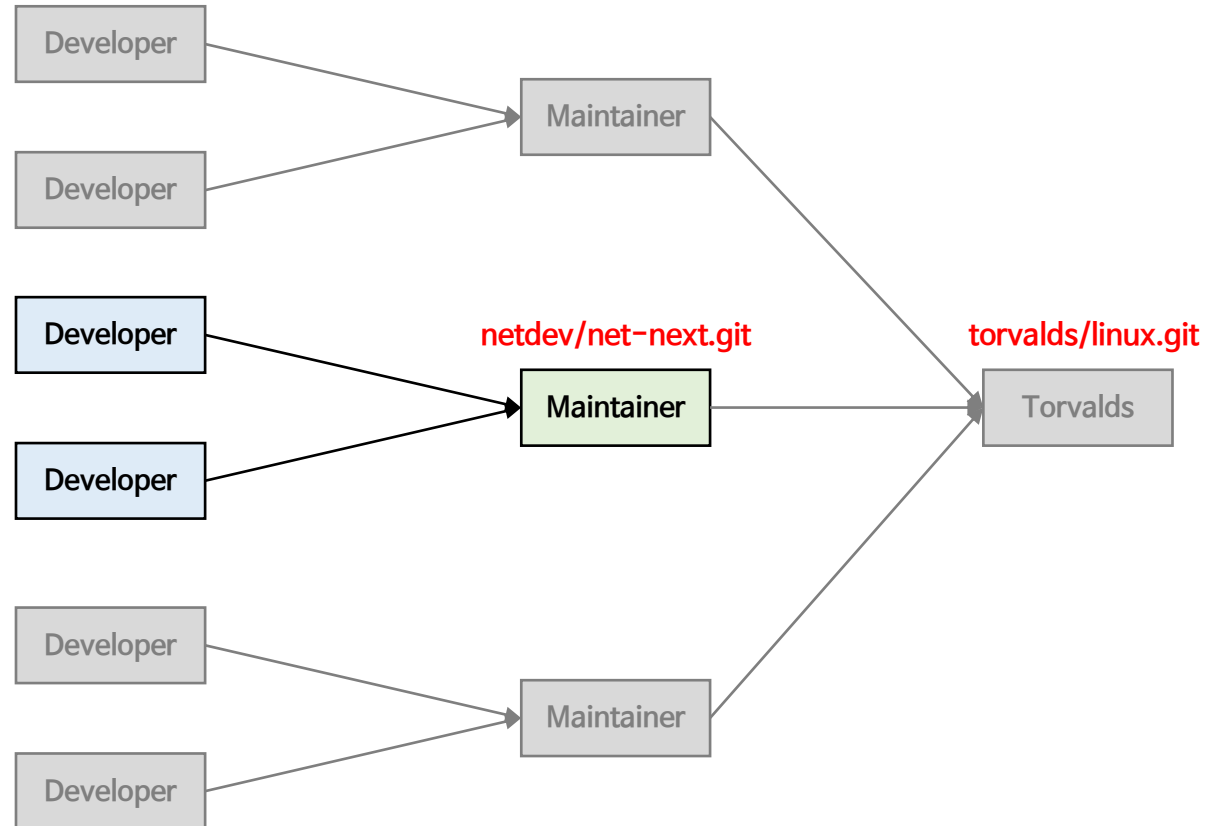


index : kernel/git/netdev/net-next.git

git.kernel.org/pub/scm/linux/kernel/git/netdev/net-next.git

Netdev Group

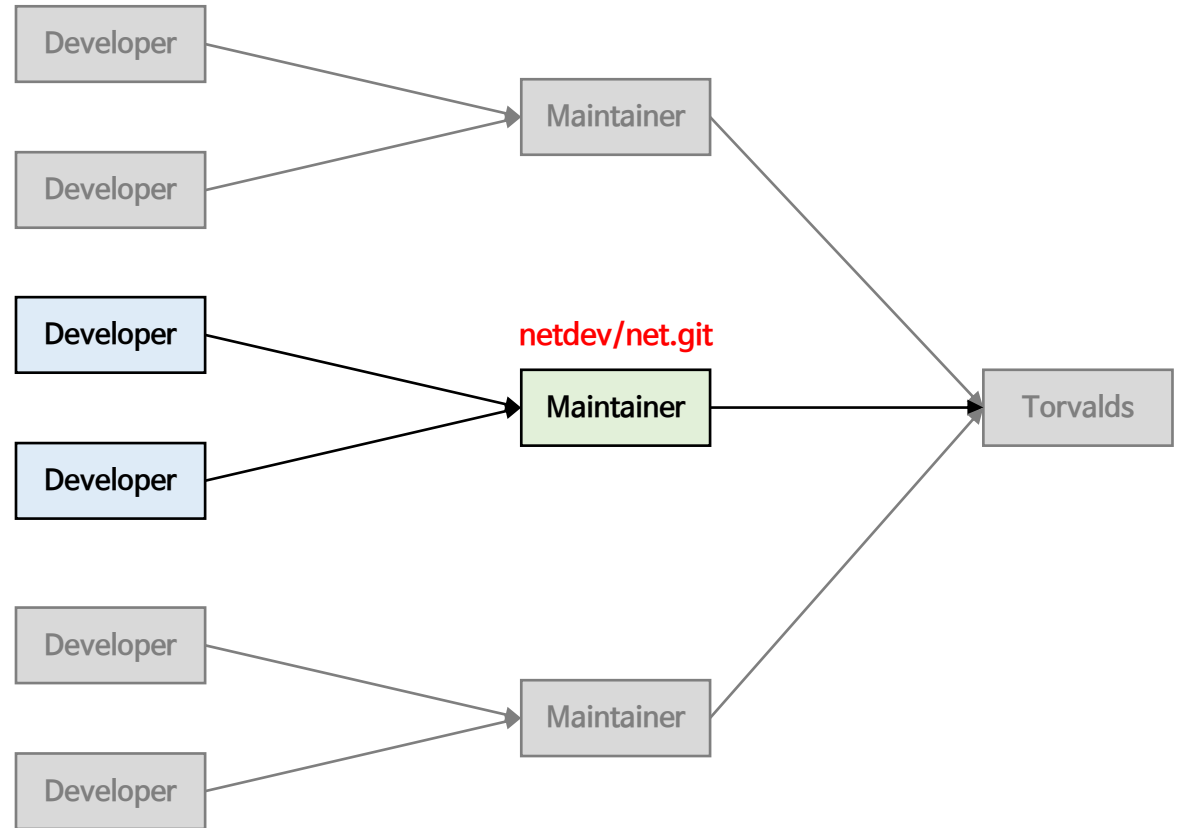
netdev/net-next.git : 기능 개선 / 버그 수정 tree



Subsystem Maintainers

- MAINTAINERS 파일

```
$ cat MAINTAINERS | grep -A 30 'NETWORKING \[GENERAL\]'  
NETWORKING [GENERAL]  
M: "David S. Miller" <davem@davemloft.net>  
M: Eric Dumazet <edumazet@google.com>  
M: Jakub Kicinski <kuba@kernel.org>  
M: Paolo Abeni <pabeni@redhat.com>  
L: netdev@vger.kernel.org  
S: Maintained  
Q: https://patchwork.kernel.org/project/netdevbpf/list/  
B: mailto:netdev@vger.kernel.org  
T: git git://git.kernel.org/.../netdev/net.git  
T: git git://git.kernel.org/.../netdev/net-next.git  
F: Documentation/networking/  
F: Documentation/process/maintainer-netdev.rst  
F: include/linux/in.h  
F: include/linux/net.h  
F: include/linux/netdevice.h  
F: include/net/  
F: include/uapi/linux/in.h  
F: include/uapi/linux/net.h  
F: include/uapi/linux/net_namespace.h  
F: include/uapi/linux/netdevice.h  
F: lib/net_utils.c  
F: lib/random32.c  
F: net/  
F: tools/testing/selftests/net/
```



Mailing List

각 서브시스템별로 mailing list를 가지며
컨트리뷰터들의 패치나 discussion이 이루어진다.

```
$ cat MAINTAINERS | grep -A 30 'NETWORKING \[GENERAL]'  
NETWORKING [GENERAL]  
M: "David S. Miller" <davem@davemloft.net>  
M: Eric Dumazet <edumazet@google.com>  
M: Jakub Kicinski <kuba@kernel.org>  
M: Paolo Abeni <pabeni@redhat.com>  
L: netdev@vger.kernel.org  
S: Maintained
```

Majordomo lists at VGER.KERNEL.ORG

REMEMBER: Subscription to these lists go via <majordomo@vger.kernel.org> !

Note about archives: Listed archives are those that have been reported to vger's maintainers, or that we have found out otherwise. As things are, *list of archives is not complete.*

[alsa-devel](#), [autofs](#), [backports](#), [bpf](#), [ceph-devel](#), [cgroups](#), [cpufreq](#), [dash](#), [dccp](#), [devicetree-compiler](#), [devicetree-spec](#), [devicetree](#), [dmaengine](#), [dwarves](#), [ecryptfs](#), [fio](#), [fstests](#), [git-commits-24](#), [git-commits-head](#), [git](#), [hail-devel](#), [initramfs](#), [io-uring](#), [irda-users](#), [kernel-janitors](#), [kernel-packagers](#), [kernel-testers](#), [keyrings](#), [kvm-commits](#), [kvm-ia64](#), [kvm-ppc](#), [kvm](#), [lartc](#), [libzbc](#), [linux-8086](#), [linux-acpi](#), [linux-admin](#), [linux-alpha](#), [linux-api](#), [linux-apps](#), [linux-arch](#), [linux-arm-msm](#), [linux-assembly](#), [linux-bbs](#), [linux-bcache](#), [linux-bcachefs](#), [linux-block](#), [linux-bluetooth](#), [linux-btrace](#), [linux-btrfs](#), [linux-c-programming](#), [linux-can](#), [linux-cifs](#), [linux-clk](#), [linux-config](#), [linux-console](#), [linux-coverity](#), [linux-crypto](#), [linux-csky](#), [linux-cxl](#), [linux-diald](#), [linux-doc](#), [linux-edac](#), [linux-efi](#), [linux-embedded](#), [linux-ext4](#), [linux-fbdev](#), [linux-fido](#), [linux-fpga](#), [linux-fscrypt](#), [linux-fsdevel](#), [linux-fsf](#), [linux-ftp](#), [linux-gcc](#), [linux-gpio](#), [linux-hams](#), [linux-hardening](#), [linux-hexagon](#), [linux-hotplug](#), [linux-hwmon](#), [linux-hyperv](#), [linux-i2c](#), [linux-ia64](#), [linux-ibcs2](#), [linux-ide](#), [linux-iiio](#), [linux-input](#), [linux-integrity](#), [linux-ipx](#), [linux-isdn](#), [linux-japanese](#), [linux-kbuild](#), [linux-kernel-announce](#), [linux-kernel-announce-posters](#), [linux-kernel](#), [linux-kseltest](#), [linux-laptop](#), [linux-leds](#), [linux-linuxss](#), [linux-lvm](#), [linux-lugnuts](#), [linux-m68k-cvscommit](#), [linux-m68k](#), [linux-man](#), [linux-mca](#), [linux-media](#), [linux-metag](#), [linux-mips](#), [linux-mmc](#), [linux-modules](#), [linux-msdos-devel](#), [linux-msdos](#), [linux-new-lists](#), [linux-newbie](#), [linux-next](#), [linux-nfs](#), [linux-nilfs](#), [linux-numa](#), [linux-omap](#), [linux-opengl](#), [linux-parisc](#), [linux-pci](#), [linux-perf-users](#), [linux-pm](#), [linux-ppp](#), [linux-pwm](#), [linux-raid](#), [linux-rdma](#), [linux-remoteproc](#), [linux-renesas-soc](#), [linux-rt-users](#), [linux-rtc](#), [linux-s390](#), [linux-samsung-soc](#), [linux-scsi](#), [linux-sctp](#), [linux-security-module](#), [linux-serial](#), [linux-sgx](#), [linux-sh](#), [linux-smp](#), [linux-soc](#), [linux-sound](#), [linux-sparse](#), [linux-spdx](#), [linux-spi](#), [linux-standards](#), [linux-svgalib](#), [linux-tape](#), [linux-tegra](#), [linux-tip-commits](#), [linux-toolchains](#), [linux-trace-devel](#), [linux-trace-kernel](#), [linux-trace-users](#), [linux-unionfs](#), [linux-usb](#), [linux-userfs](#), [linux-watchdog](#), [linux-wireless](#), [linux-word](#), [linux-wpan](#), [linux-x11](#), [linux-x25](#), [linux-x86_64](#), [linux-xfs](#), [live-patching](#), [lm-sensors](#), [lpc-bpf](#), [lpc-netdev-bpf](#), [lpc-netdev](#), [lvs-devel](#), [mm-commits](#), [netconf2019](#), [netdev-driver-reviewers](#), [netdev](#), [netfilter-devel](#), [netfilter](#), [perfbook](#), [phone-devel](#), [platform-driver-x86](#), [rcu](#), [reiserfs-devel](#), [rust-for-linux](#), [selinux-refpolicy](#), [selinux](#), [smatch](#), [sparclinux](#), [stable-commits](#), [stable-rt](#), [stable](#), [stgt](#), [target-devel](#), [testing](#), [trinity](#), [ultraio](#), [util-linux](#), [workflows](#), [xdp-newbies](#),

Listname	Number of subscribers
List: alsa-devel ; (subscribe / unsubscribe) Info: Archives: http://www.spinics.net/lists/alsa-devel/ Footer: --- To unsubscribe from this list: send the line "unsubscribe alsa-devel" in	252
List: autofs ; (subscribe / unsubscribe) Info: This is the mailing list for autofs development. Archives: Footer:	108

<http://vger.kernel.org/vger-lists.html>


Mailing List

[PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open code

Open code which is dev->priv_flags & IFF_MACSEC has already defined as netif_is_macsec(). So use netif_is_macsec() instead of open code. This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>

컨트리뷰터가 패치 전송




Mailing List

```
[PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open code
```

Open code which is dev->priv_flags & IFF_MACSEC has already defined as netif_is_macsec(). So use netif_is_macsec() instead of open code. This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>

Mailing List에서 Review 와 Discussion 발생




Mailing List

패치가 **accepted** 되면 서브시스템에 반영되었다는 메일이 전송됨

```
[PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open code
```

Open code which is dev->priv_flags & IFF_MACSEC has already defined as netif_is_macsec(). So use netif_is_macsec() instead of open code. This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>



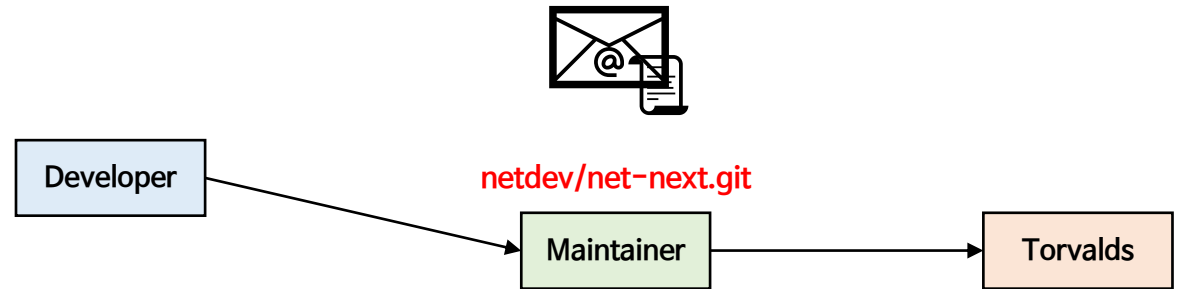
```
Re: [PATCH net-next] mlxsw: spectrum: use netif_is_macsec() instead of open ode
```

Hello:

This series was applied to netdev/net-next.git (master) by David S. Miller <davem@davemloft.net>:

Subsystem Maintainers

패치는 net-next git tree에 반영



Subsystem Maintainers



index : kernel/git/torvalds/linux.git

Linux kernel source tree

about summary refs log tree **commit** diff stats

author Linus Torvalds <torvalds@linux-foundation.org> 2022-03-17 12:55:26 -0700
committer Linus Torvalds <torvalds@linux-foundation.org> 2022-03-17 12:55:26 -0700
commit 551acdc3c3d2b6bc97f11e31dcf960bc36343bfc (patch)
tree 1a2bf01b80da58fc613735b1c3fb03e5fb55bc2c
parent c81801eb7f2476a25d8fb27449e01b0bef46908a (diff)
parent b04683ff8f0823b869c219c78ba0d974bddea0b5 (diff)
download linux-551acdc3c3d2b6bc97f11e31dcf960bc36343bfc.tar.gz


Merge tag 'net-5.17-final' of git://git.kernel.org/pub/scm/linux/kernel/git/netdev/net

Pull networking fixes from Jakub Kicinski:
"Including fixes from netfilter, ipsec, and wireless.

A few last minute revert / disable and fix patches came down from our sub-trees. We're not waiting for any fixes at this point.

Current release - regressions:

- Revert "netfilter: nat: force port remap to prevent shadowing well-known ports", restore working conntrack on asymmetric paths
- Revert "ath10k: drop beacon and probe response which leak from other channel", restore working AP and mesh mode on QCA9984
- eth: intel: fix hang during reboot/shutdown



상위 tree로 Pull Request 전송

Subsystem Maintainers



index : kernel/git/torvalds/linux.git

Linux kernel source tree

about summary refs log tree **commit** diff stats

```
author Linus Torvalds <torvalds@linux-foundation.org> 2022-03-17 12:55:26 -0700
committer Linus Torvalds <torvalds@linux-foundation.org> 2022-03-17 12:55:26 -0700
commit 551acdc3c3d2b6bc97f11e31dcf960bc36343bfc (patch)
tree 1a2bf01b80da58fc613735b1c3fb03e5fb55bc2c
parent c81801eb7f2476a25d8fb27449e01b0bef46908a (diff)
parent b04683ff8f0823b869c219c78ba0d974bddea0b5 (diff)
download linux-551acdc3c3d2b6bc97f11e31dcf960bc36343bfc.tar.gz
```

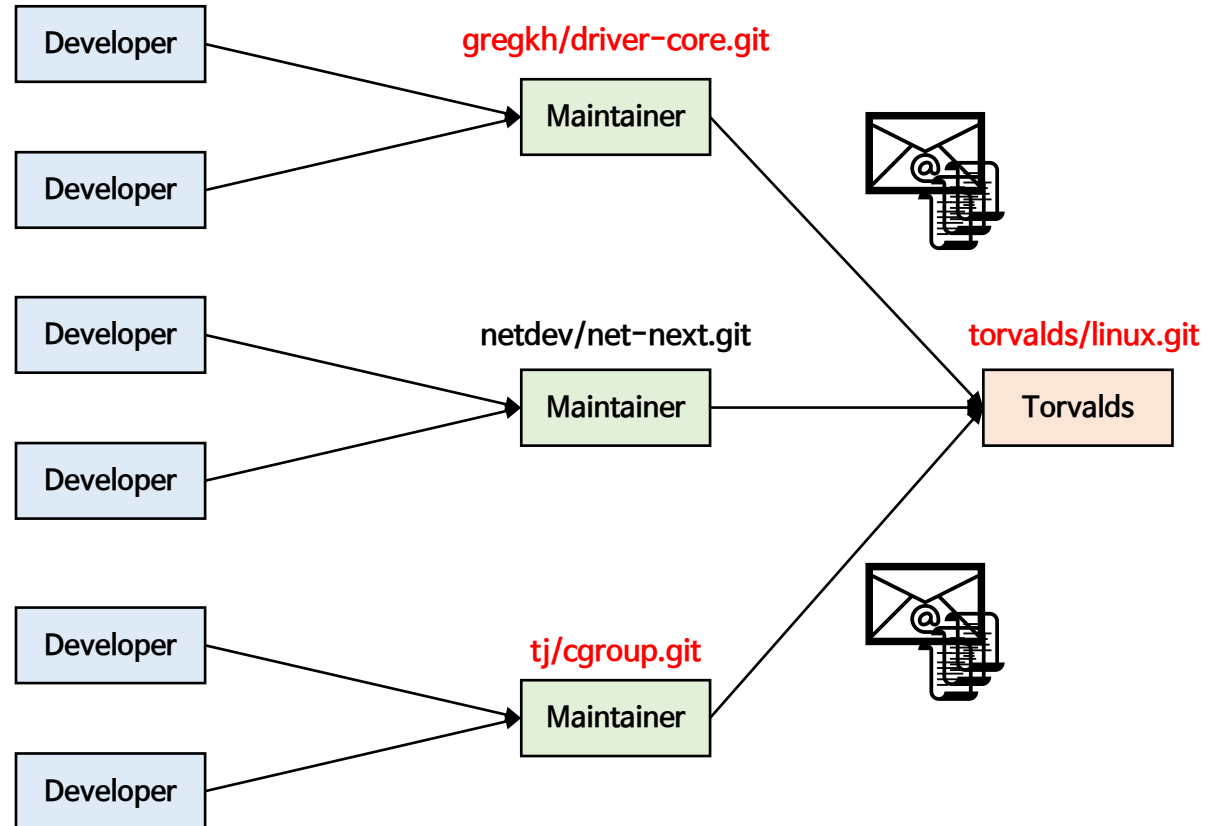
Merge tag 'net-5.17-final' of git://git.kernel.org/pub/scm/linux/kernel/git/netdev/net

Pull networking fixes from Jakub Kicinski:
"Including fixes from netfilter, ipsec, and wireless."

A few last minute revert / disable and fix patches came down from our sub-trees. We're not waiting for any fixes at this point.

Current release - regressions:

- Revert "netfilter: nat: force port remap to prevent shadowing well-known ports", restore working conntrack on asymmetric paths
- Revert "ath10k: drop beacon and probe response which leak from other channel", restore working AP and mesh mode on QCA9984
- eth: intel: fix hang during reboot/shutdown



다양한 서브시스템에서도 상위 tree로 Pull Request 전송

Patch

리눅스 커널의 기여 단위

리눅스 커널의 기여 단위

개발자들은 개발 내용을 patch로 만들어서
이메일로 메일링 리스트에 제출

리눅스 커널의 기여 단위



개발자들은 개발 내용을 patch로 만들어서
이메일로 메일링 리스트에 제출

```
$ git format-patch -1  
--subject-prefix="PATCH, net-next"
```

리눅스 커널의 기여 단위



개발자들은 개발 내용을 patch로 만들어서
이메일로 메일링 리스트에 제출

```
$ git format-patch -1  
--subject-prefix="PATCH, net-next"
```



```
From 0199215216978b612d4e8be11e878b87bc643033 Mon Sep 17 00:00:00 2001  
From: Juhee Kang <claudiajkang@gmail.com>  
Date: Sun, 10 Oct 2021 13:03:29 +0900  
Subject: [PATCH, net-next] mlxsw: spectrum: use netif_is_macsec() instead of  
open code  
  
Open code which is dev->priv_flags & IFF_MACSEC has already defined as  
netif_is_macsec(). So use netif_is_macsec() instead of open code.  
This patch doesn't change logic.  
  
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>  
---  
include/linux/netdevice.h | 2 +-  
1 file changed, 1 insertion(+), 1 deletion(-)  
  
diff --git a/include/linux/netdevice.h b/include/linux/netdevice.h  
index 15f4a658e436..0723c1314ea2 100644  
--- a/include/linux/netdevice.h  
+++ b/include/linux/netdevice.h  
@@ -5237,7 +5237,7 @@ static inline void netif_keep_dst(struct net_device *dev)  
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)  
{  
    /* TODO: reserve and use an additional IFF bit, if we get more users */  
-    return dev->priv_flags & IFF_MACSEC;  
+    return netif_is_macsec(dev);  
}  
  
extern struct pernet_operations __net_initdata loopback_net_ops;  
--  
2.34.1
```

리눅스 커널의 기여 단위



개발자들은 개발 내용을 patch로 만들어서
이메일로 메일링 리스트에 제출



```
$ git format-patch -1  
--subject-prefix="PATCH, net-next"
```

```
$ git send-patch --cc "Network Maintainer  
<example@email.com>" --to "Network  
Mailing List <list@email.com>" *.patch
```

리눅스 커널의 기여 단위

개발자들은 개발 내용을 patch로 만들어서
이메일로 메일링 리스트에 제출



```
$ git format-patch -1  
--subject-prefix="PATCH, net-next"
```

```
$ git send-patch --cc "Network Maintainer  
<example@email.com>" --to "Network  
Mailing List <list@email.com>" *.patch
```



Patchwork Netdev + BPF Patches Bundles About this Login Register Mail settings

Show patches with: Submitter = Juhee Kang | 29 patches

Patch	Date	Submitter	Delegate	State
[net-next,3/3] pktgen: document the latest pktgen usage options	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/3] samples: pktgen: add trap SIGINT for printing execution result	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,1/3] samples: pktgen: fix to print when terminated normally	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/2] samples: pktgen: add missing IPv6 option to pktgen scripts	2021-08-12	Juhee Kang	netdev	Accepted
[net-next,1/2] samples: pktgen: pass the environment variable of normal user to suc	2021-08-12	Juhee Kang	netdev	Accepted
[bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-24	Juhee Kang	bpf	Superseded
[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-24	Juhee Kang	bpf	Superseded


리눅스 커널의 기여 단위

개발자들은 개발 내용을 patch로 만들어서
이메일로 메일링 리스트에 제출



```
$ git format-patch -1  
--subject-prefix="PATCH, net-next"
```

```
$ git send-patch --cc "Network Maintainer  
<example@email.com>" --to "Network  
Mailing List <list@email.com>" *.patch
```



Patchwork Netdev + BPF Patches Bundles About this Login Register Mail settings

Show patches with: Submitter = Juhee Kang | 29 patches

Patch	Date	Submitter	Delegate	State
[net-next,3/3] pktgen: document the latest pktgen usage options	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/3] samples: pktgen: add trap SIGINT for printing execution result	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,1/3] samples: pktgen: fix to print when terminated normally	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/2] samples: pktgen: add missing IPv6 option to pktgen scripts	2021-08-12	Juhee Kang	netdev	Accepted
[net-next,1/2] samples: pktgen: pass the environment variable of normal user to suc	2021-08-12	Juhee Kang	netdev	Accepted
[bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-24	Juhee Kang	bpf	Superseded
[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-24	Juhee Kang	bpf	Superseded

Patch 구성

```
commit 0199215216978b612d4e8be11e878b87bc643033
Author: Juhee Kang <claudiajkang@gmail.com>
AuthorDate: Sun Oct 10 13:03:29 2021 +0900
Commit: David S. Miller <davem@davemloft.net>
CommitDate: Sun Oct 10 11:18:48 2021 +0100
```

```
mlxsw: spectrum: use netif_is_macsec() instead of open code
```

```
Open code which is dev->priv_flags & IFF_MACSEC has already defined as
netif_is_macsec(). So use netif_is_macsec() instead of open code.
This patch doesn't change logic.
```

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
```

```
---
include/linux/netdevice.h | 2 +-
1 file changed, 1 insertion(+), 1 deletion(-)
```

```
diff --git a/include/linux/netdevice.h b/include/linux/netdevice.h
index 15f4a658e436..0723c1314ea2 100644
--- a/include/linux/netdevice.h
+++ b/include/linux/netdevice.h
@@ -5237,7 +5237,7 @@ static inline void netif_keep_dst(struct net_device *dev)
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)
{
    /* TODO: reserve and use an additional IFF bit, if we get more users */
-   return dev->priv_flags & IFF_MACSEC;
+   return netif_is_macsec(dev);
}

extern struct pernet_operations __net_initdata loopback_net_ops;
--
2.34.1
```

Commit ID

commit을 중복 없이 기억하기 위해
SHA1 알고리즘으로 만들어진 hash 값을 사용

Patch 구성

```
commit 0199215216978b612d4e8be11e878b87bc643033
Author: Juhee Kang <claudiajkang@gmail.com>
AuthorDate: Sun Oct 10 13:03:29 2021 +0900
Commit: David S. Miller <davem@davemloft.net>
CommitDate: Sun Oct 10 11:18:48 2021 +0100
```

```
mlxsw: spectrum: use netif_is_macsec() instead of open code
```

```
Open code which is dev->priv_flags & IFF_MACSEC has already defined as
netif_is_macsec(). So use netif_is_macsec() instead of open code.
This patch doesn't change logic.
```

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
```

```
---
include/linux/netdevice.h | 2 +-
1 file changed, 1 insertion(+), 1 deletion(-)

diff --git a/include/linux/netdevice.h b/include/linux/netdevice.h
index 15f4a658e436..0723c1314ea2 100644
--- a/include/linux/netdevice.h
+++ b/include/linux/netdevice.h
@@ -5237,7 +5237,7 @@ static inline void netif_keep_dst(struct net_device *dev)
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)
{
    /* TODO: reserve and use an additional IFF bit, if we get more users */
-   return dev->priv_flags & IFF_MACSEC;
+   return netif_is_macsec(dev);
}

extern struct pernet_operations __net_initdata loopback_net_ops;
--
2.34.1
```

Commit ID

Author

패치를 작성한 사람의 정보
ex) 작성자 이름 <작성자 메일 정보>

Author Date

작성자가 해당 커밋을 작성한 시간

Patch 구성

```
commit 0199215216978b612d4e8be11e878b87bc643033
Author: Juhee Kang <claudiajkang@gmail.com>
AuthorDate: Sun Oct 10 13:03:29 2021 +0900
Commit: David S. Miller <davem@davemloft.net>
CommitDate: Sun Oct 10 11:18:48 2021 +0100
```

```
mlxsw: spectrum: use netif_is_macsec() instead of open code
```

```
Open code which is dev->priv_flags & IFF_MACSEC has already defined as
netif_is_macsec(). So use netif_is_macsec() instead of open code.
This patch doesn't change logic.
```

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
```

```
---
```

```
include/linux/netdevice.h | 2 +-
1 file changed, 1 insertion(+), 1 deletion(-)
```

```
diff --git a/include/linux/netdevice.h b/include/linux/netdevice.h
index 15f4a658e436..0723c1314ea2 100644
--- a/include/linux/netdevice.h
+++ b/include/linux/netdevice.h
@@ -5237,7 +5237,7 @@ static inline void netif_keep_dst(struct net_device *dev)
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)
{
    /* TODO: reserve and use an additional IFF bit, if we get more users */
-   return dev->priv_flags & IFF_MACSEC;
+   return netif_is_macsec(dev);
}
```

```
extern struct pernet_operations __net_initdata loopback_net_ops;
--
2.34.1
```

Commit ID

Author

Author Date

Commit

리눅스 커널에서 committer은 해당 패치를
git 저장소에 반영한 사람을 의미

ex) Maintainer 혹은 Reviewer

Commit Date

해당 커밋이 반영된 시간

Patch 구성

```
commit 0199215216978b612d4e8be11e878b87bc643033
Author: Juhee Kang <claudiajkang@gmail.com>
AuthorDate: Sun Oct 10 13:03:29 2021 +0900
Commit: David S. Miller <davem@davemloft.net>
CommitDate: Sun Oct 10 11:18:48 2021 +0100

    mlxsw: spectrum: use netif_is_macsec() instead of open code

Open code which is dev->priv_flags & IFF_MACSEC has already defined as
netif_is_macsec(). So use netif_is_macsec() instead of open code.
This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
---
include/linux/netdevice.h | 2 +-
1 file changed, 1 insertion(+), 1 deletion(-)

diff --git a/include/linux/netdevice.h b/include/linux/netdevice.h
index 15f4a658e436..0723c1314ea2 100644
--- a/include/linux/netdevice.h
+++ b/include/linux/netdevice.h
@@ -5237,7 +5237,7 @@ static inline void netif_keep_dst(struct net_device *dev)
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)
{
    /* TODO: reserve and use an additional IFF bit, if we get more users */
-   return dev->priv_flags & IFF_MACSEC;
+   return netif_is_macsec(dev);
}

extern struct pernet_operations __net_initdata loopback_net_ops;
--
2.34.1
```

Commit ID
Author
Author Date
Commit
Commit Date
Commit Header

해당 커밋에 대한 간략한 설명
다음과 같은 형식으로 작성

ex) **모듈 이름: 세부 모듈: 커밋 설명**

Commit Description

해당 커밋에 대한 상세한 설명을 작성

Patch 구성

```
commit 0199215216978b612d4e8be11e878b87bc643033
Author: Juhee Kang <claudiajkang@gmail.com>
AuthorDate: Sun Oct 10 13:03:29 2021 +0900
Commit: David S. Miller <davem@davemloft.net>
CommitDate: Sun Oct 10 11:18:48 2021 +0100

    mlxsw: spectrum: use netif_is_macsec() instead of open code

Open code which is dev->priv_flags & IFF_MACSEC has already defined as
netif_is_macsec(). So use netif_is_macsec() instead of open code.
This patch doesn't change logic.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
---
include/linux/netdevice.h | 2 +-
1 file changed, 1 insertion(+), 1 deletion(-)

diff --git a/include/linux/netdevice.h b/include/linux/netdevice.h
index 15f4a658e436..0723c1314ea2 100644
--- a/include/linux/netdevice.h
+++ b/include/linux/netdevice.h
@@ -5237,7 +5237,7 @@ static inline void netif_keep_dst(struct net_device *dev)
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)
{
    /* TODO: reserve and use an additional IFF bit, if we get more users */
-   return dev->priv_flags & IFF_MACSEC;
+   return netif_is_macsec(dev);
}

extern struct pernet_operations __net_initdata loopback_net_ops;
--
2.34.1
```

Commit ID
Author
Author Date
Commit
Commit Date
Commit Header
Commit Description
Commit Tag

Signed-off-by:

작성자는 라이선싱 정책에 대해 이해했고
커밋을 오픈소스 패치로 제출하는 것에 동의

ex) **Signed-off-by:** 이름 <메일>

Patch 구성

```
mlxsw: spectrum: use netif_is_macsec() instead of open code
=== Commit Body ===
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```

```
samples: bpf: Fix tracex7 error raised on the missing argument
=== Commit Body ===
Acked-by: Yonghong Song <yhs@fb.com>
```

```
amt: Use BIT macros instead of open codes
=== Commit Body ===
Reviewed-by: Taehee Yoo <ap420073@gmail.com>
Tested-by: Taehee Yoo <ap420073@gmail.com>
```

```
net: hsr: fix hsr build error when lockdep is not enabled
=== Commit Body ===
Reported-by: Eric Dumazet <eric.dumazet@gmail.com>
```

```
bpf_glue: include errno.h
=== Commit Body ===
Fixes: ac4e0913beb1 ("bpf: Export bpf syscall wrapper")
```

Commit Tag Detail

1) Signed-off-by

라이센싱 정책에 대해 이해했고, 커밋을 오픈소스 패치로 제출하는 것에 동의

2) Acked-by

메인테이너가 해당 패치를 확인하였고, 개발자가 작성한 패치를 반영하기로 결정

3) Reviewed-by

해당 패치의 코드에 대해 Reviewer 혹은 다른 개발자가 리뷰하였음

4) Tested-by

해당 패치의 변경사항에 대해 Reviewer 혹은 다른 개발자가 테스트하였음

5) Reported-by / Suggested-by

해당 버그를 발견하거나 기능을 제안한 사람을 명시하기 위해 남기는 태그

6) Fixes

해당 패치가 어떤 커밋에서 발생한 버그를 해결하는지를 추적하기 위한 태그

Patch 구성

```
commit 0199215216978b612d4e8be11e878b87bc643033
Author: Juhee Kang <claudiajkang@gmail.com>
AuthorDate: Sun Oct 10 13:03:29 2021 +0900
Commit: David S. Miller <davem@davemloft.net>
CommitDate: Sun Oct 10 11:18:48 2021 +0100
```

```
mlxsw: spectrum: use netif_is_macsec() instead of open code
```

```
Open code which is dev->priv_flags & IFF_MACSEC has already defined as
netif_is_macsec(). So use netif_is_macsec() instead of open code.
This patch doesn't change logic.
```

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>
```

```
---
include/linux/netdevice.h | 2 +-
1 file changed, 1 insertion(+), 1 deletion(-)

diff --git a/include/linux/netdevice.h b/include/linux/netdevice.h
index 15f4a658e436..0723c1314ea2 100644
--- a/include/linux/netdevice.h
+++ b/include/linux/netdevice.h
@@ -5237,7 +5237,7 @@ static inline void netif_keep_dst(struct net_device *dev)
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)
{
    /* TODO: reserve and use an additional IFF bit, if we get more users */
-   return dev->priv_flags & IFF_MACSEC;
+   return netif_is_macsec(dev);
}

extern struct pernet_operations __net_initdata loopback_net_ops;
--
2.34.1
```

Commit ID
Author
Author Date
Commit
Commit Date
Commit Header
Commit Description
Commit Tag
diff 결과 (스냅샷 정보)

해당 커밋에 포함된 diff 결과

- 수정이 발생한 파일 목록

```
include/linux/netdevice.h | 2 +-
1 file changed, 1 insertion(+), 1 deletion(-)
```

- 세부 수정사항

```
static inline bool netif_reduces_vlan_mtu(struct net_device *dev)
{
    /* TODO: reserve and use an additional IFF bit, if we get more
users */
-   return dev->priv_flags & IFF_MACSEC;
+   return netif_is_macsec(dev);
```

리눅스 Kernel에 이제 실제로 Contribution 해볼까요?

3

Kernel 기여하기

컨트리뷰션 방법

- ① 문서 기여
- ② 버그 수정
- ③ 기능 개선

1. 문서 기여



index : kernel/git/next/linux-next.git

The linux-next integration testing tree

[about](#) [summary](#) [refs](#) [log](#) [tree](#) **[commit](#)** [diff](#) [stats](#)

```
author Juhee Kang <claudiajkang@gmail.com> 2021-08-25 19:57:17 +0900
committer David S. Miller <davem@davemloft.net> 2021-08-25 13:44:30 +0100
commit 246b184fffdcead3710228e3bff744ce8c9828a3 (patch)
tree be72f30a9f1135e1c8cb1524f8f031bc5f9a0f0e
parent 6c882bdc4bcd63e164f05738e7677b8a62fc0ec1 (diff)
download linux-next-246b184fffdcead3710228e3bff744ce8c9828a3.tar.gz
```

pktgen: document the latest pktgen usage options

Currently, the pktgen.rst documentation doesn't cover the latest pktgen sample usage options such as count and IPv6, and so on. Also, this documentation includes the old sample scripts which are no longer use because it was removed by the commit a4b6ade8359f ("samples/pktgen : remove remaining old pktgen sample scripts")

Thus, this commit documents pktgen sample usage using the latest options and removes old sample scripts, and fixes a minor typo.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>

Diffstat

```
-rw-r--r-- Documentation/networking/pktgen.rst 18
```

1 files changed, 8 insertions, 10 deletions

```
diff --git a/Documentation/networking/pktgen.rst b/Documentation/networking/pktgen.rst
```

```
index 7afac9f11838..1225f0f63f07 100644
```

```
--- a/Documentation/networking/pktgen.rst
```

```
+++ b/Documentation/networking/pktgen.rst
```

```
@@ -248,26 +248,24 @@ Usage:::
```

```
-i : ($DEV)          output interface/device (required)
-s : ($PKT_SIZE)    packet size
- -d : ($DEST_IP)   destination IP
+ -d : ($DEST_IP)   destination IP. CIDR (e.g. 198.18.0.0/15) is also allowed
-m : ($DST_MAC)    destination MAC-addr
+ -p : ($DST_PORT) destination PORT range (e.g. 433-444) is also allowed
-t : ($THREADS)    threads to start
+ -f : ($F_THREAD)  index of first thread (zero indexed CPU number)
-c : ($SKB_CLONE)  SKB clones send before alloc new SKB
+ -n : ($COUNT)   num messages to send per thread, 0 means indefinitely
-b : ($BURST)     HW level bursting of SKBs
-v : ($VERBOSE)   verbose
-x : ($DEBUG)     debug
+ -6 : ($IPv6)     IPv6
+ -w : ($DELAY)    Tx Delay value (ns)
+ -a : ($APPEND)   Script will not reset generator's state, but will append its config
```

신규 문서 추가

기존 문서 수정

2. 버그 수정



index : kernel/git/next/linux-next.git

The linux-next integration testing tree

[about](#) [summary](#) [refs](#) [log](#) [tree](#) [commit](#) [diff](#) [stats](#)

```
author    Juhee Kang <claudiajkang@gmail.com> 2021-10-04 21:14:38 +0900
committer Pablo Neira Ayuso <pablo@netfilter.org> 2021-10-07 19:35:57 +0200
commit    902c0b1887522a099aa4e1e6b4b476c2fe5dd13e (patch)
tree      37a3ee3d04c0a7d451b3d6c30f1794bb6d266276
parent    8d6c414cd2fb74aa6812e9bfec6178f8246c4f3a (diff)
download  linux-next-902c0b1887522a099aa4e1e6b4b476c2fe5dd13e.tar.gz
```

netfilter: xt_IDLETIMER: fix panic that occurs when timer_type has garbage value

Currently, when the rule related to IDLETIMER is added, idletimer_tg timer structure is initialized by kcalloc on executing idletimer_tg_create function. However, in this process timer->timer_type is not defined to a specific value. Thus, timer->timer_type has garbage value and it occurs kernel panic. So, this commit fixes the panic by initializing timer->timer_type using kzalloc instead of kcalloc.

```
Test commands:
# iptables -A OUTPUT -j IDLETIMER --timeout 1 --label test
$ cat /sys/class/xt_idletimer/timers/test
Killed
```

```
Splat looks like:
BUG: KASAN: user-memory-access in alarm_expires_remaining+0x49/0x70
Read of size 8 at addr 0000002e8c7bc4c8 by task cat/917
CPU: 12 PID: 917 Comm: cat Not tainted 5.14.0+ #3 79940a339f71eb14fc81aee1757a20d5bf13eb0e
Hardware name: QEMU Standard PC (Q35 + ICH9, 2009), BIOS 1.13.0-lubuntu1.1 04/01/2014
Call Trace:
 dump_stack_lvl+0x6e/0x9c
 kasan_report.cold+0x112/0x117
 ? alarm_expires_remaining+0x49/0x70
 __asan_load8+0x86/0xb0
 alarm_expires_remaining+0x49/0x70
 idletimer_tg_show+0xe5/0x19b [xt_IDLETIMER 11219304af9316a21bee5ba9d58f76a6b9bcc6d]
 dev_attr_show+0x3c/0x60
 sysfs_kf_seq_show+0x11d/0x1f0
 ? device_remove_bin_file+0x20/0x20
 kernfs_seq_show+0xa4/0xb0
```

버그가 잘 동작하도록 수정 후 패치 전송

3. 기능 개선



index : kernel/git/next/linux-next.git

The linux-next integration testing tree

about summary refs log tree **commit** diff stats

```

author       Juhee Kang <claudiajkang@gmail.com> 2021-08-13 00:08:13 +0900
committer    David S. Miller <davem@davemloft.net> 2021-08-16 11:02:09 +0100
commit       0f0c4f1b72e090b23131700bb155944cc28b2a7b (patch)
tree         1dc881a9d04a4256e1f91415a6e7f9ec0796b729
parent       7caeabd726f22e6a6c44c434574fb489986e5baa (diff)
download     linux-next-0f0c4f1b72e090b23131700bb155944cc28b2a7b.tar.gz

```

samples: pktgen: add missing IPv6 option to pktgen scripts

Currently, "sample04" and "sample05" are not working properly when running with an IPv6 option ("-6"). The commit 0f06a6787e05 ("samples: Add an IPv6 "-6" option to the pktgen scripts") has omitted the addition of this option at "sample04" and "sample05".

In order to support IPv6 option, this commit adds logic related to IPv6 option.

Fixes: 0f06a6787e05 ("samples: Add an IPv6 "-6" option to the pktgen scripts")

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>

Diffstat

```

-rwxr-xr-x samples/pktgen/pktgen_sample04_many_flows.sh 12
-rwxr-xr-x samples/pktgen/pktgen_sample05_flow_per_thread.sh 12

```

2 files changed, 14 insertions, 10 deletions

```

diff --git a/samples/pktgen/pktgen_sample04_many_flows.sh b/samples/pktgen/pktgen_sample04_many_flows.sh
index 56c5f5af350f6..cf151f861506d 100755
--- a/samples/pktgen/pktgen_sample04_many_flows.sh
+++ b/samples/pktgen/pktgen_sample04_many_flows.sh
@@ -13,13 +13,15 @@ root_check_run_with_sudo "$@"
 # Parameter parsing via include
 source ${basedir}/parameters.sh
 # Set some default params, if they didn't get set
-[ -z "$DEST_IP" ]    && DEST_IP="198.18.0.42"
+if [ -z "$DEST_IP" ]; then
+  [ -z "$IP6" ] && DEST_IP="198.18.0.42" || DEST_IP="F000::1"
+fi
[ -z "$DST_MAC" ]    && DST_MAC="90:e2:ba:ff:ff:ff"
[ -z "$CLONE_SKB" ]  && CLONE_SKB="0"
[ -z "$COUNT" ]    && COUNT="0" # Zero means indefinitely
if [ -n "$DEST_IP" ]; then
-  validate_addr $DEST_IP
-  read -r DST_MIN DST_MAX <<< $(parse_addr $DEST_IP)
+  validate_addr ${IP6} $DEST_IP
+  read -r DST_MIN DST_MAX <<< $(parse_addr ${IP6} $DEST_IP)

```

신규 기능 추가

3. 기능 개선



index : kernel/git/next/linux-next.git

The linux-next integration testing tree

about summary refs log tree **commit** diff stats

```

author Juhee Kang <claudiajkang@gmail.com> 2022-02-05 15:40:38 +0000
committer David S. Miller <davem@davemloft.net> 2022-02-06 10:55:52 +0000
commit 4acc45db711586dad98e681ccb93f29372907a2 (patch)
tree f69f105c0d756fd7caacec85ef42c2d9463843bf
parent 5a8fb33e530512ee67a11b30f3451a4f030f4b01 (diff)
download linux-next-4acc45db711586dad98e681ccb93f29372907a2.tar.gz

```

net: hsr: use hlist_head instead of list_head for mac addresses

Currently, HSR manages mac addresses of known HSR nodes by using list_head. It takes a lot of time when there are a lot of registered nodes due to finding specific mac address nodes by using linear search. We can be reducing the time by using hlist. Thus, this patch moves list_head to hlist_head for mac addresses and this allows for further improvement of network performance.

```

Condition: registered 10,000 known HSR nodes
Before:
# iperf3 -c 192.168.10.1 -i 1 -t 10
Connecting to host 192.168.10.1, port 5201
[ 5] local 192.168.10.2 port 59442 connected to 192.168.10.1 port 5201
[ ID] Interval      Transfer   Bitrate   Retr  Cwnd
[ 5] 0.00-1.49 sec   3.75 MBytes 21.1 Mbits/sec  0   158 KBytes
[ 5] 1.49-2.05 sec   1.25 MBytes 18.7 Mbits/sec  0   166 KBytes
[ 5] 2.05-3.06 sec   2.44 MBytes 20.3 Mbits/sec  56  16.9 KBytes
[ 5] 3.06-4.08 sec   1.43 MBytes 11.7 Mbits/sec  11  38.0 KBytes
[ 5] 4.08-5.00 sec    951 KBytes  8.49 Mbits/sec  0   56.3 KBytes

```

```

After:
# iperf3 -c 192.168.10.1 -i 1 -t 10
Connecting to host 192.168.10.1, port 5201
[ 5] local 192.168.10.2 port 36460 connected to 192.168.10.1 port 5201
[ ID] Interval      Transfer   Bitrate   Retr  Cwnd
[ 5] 0.00-1.00 sec   7.39 MBytes 62.0 Mbits/sec  3   130 KBytes
[ 5] 1.00-2.00 sec   5.06 MBytes 42.4 Mbits/sec  16  113 KBytes
[ 5] 2.00-3.00 sec   8.58 MBytes 72.0 Mbits/sec  42  94.3 KBytes
[ 5] 3.00-4.00 sec   7.44 MBytes 62.4 Mbits/sec  2   131 KBytes
[ 5] 4.00-5.07 sec   8.13 MBytes 63.5 Mbits/sec  38  92.9 KBytes

```

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
Signed-off-by: David S. Miller <davem@davemloft.net>

Diffstat

-rw-r--r--	net/hsr/hsr_debugfs.c	40	<div style="width: 40px; height: 10px; background-color: #28a745;"></div>
-rw-r--r--	net/hsr/hsr_device.c	10	<div style="width: 10px; height: 10px; background-color: #dc3545;"></div>
-rw-r--r--	net/hsr/hsr_forward.c	7	<div style="width: 7px; height: 10px; background-color: #dc3545;"></div>
-rw-r--r--	net/hsr/hsr_forward.h	200	<div style="width: 200px; height: 10px; background-color: #28a745;"></div>

기존 기능 강화

Kernel Patch in Action

버그를 수정하는 패치를 진행한다고 가정

문제 발견

```
$ cd samples/bpf
$ make -j
$ ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
```

커널의 테스트 스크립트
실행 도중 **에러 발생**

커널 개발 환경 구성 - 리눅스 환경 설정



Fedora
Debian
OpenSUSE
Ubuntu

가상 환경에 리눅스 설치

커널 개발 환경 구성 - 리눅스 환경 설정

```
# apt-get install build-essential vim git cscope  
libncurses-dev libssl-dev bison flex bison
```

가상 환경 리눅스에 개발 관련 패키지 설치

커널 개발 환경 구성 - 리눅스 환경 설정

```
# apt-get install build-essential vim git cscope  
libncurses-dev libssl-dev bison flex bison  
# systemctl enable serial-getty@ttyS0.service  
# systemctl start serial-getty@ttyS0.service  
# systemctl daemon-reload  
# vim /etc/default/grub          GRUB_DEFAULT=3  
                                GRUB_TIMEOUT_STYLE=menu  
                                GRUB_TIMEOUT=3  
# update-grub2
```

가상 환경 리눅스에 개발 관련 패키지 설치
및 디버깅을 위한 환경 구축

커널 개발 환경 구성 - 커널 개발 소스 코드 다운

```
$ git clone https://git.kernel.org/pub/scm/linux/kernel/git/netdev/net-next.git
```



index : kernel/git/netdev/net-next.git

Netdev Group's -next networking tree

master switch

Netdev Group

about summary refs log tree commit diff stats

log msg search

Branch	Commit message	Author	Age
main	Merge tag 'net-next-6.1' of git://git.kernel.org/pub/scm/linux/kernel/git/net...	Linus Torvalds	3 days
master	Merge tag 'net-next-6.1' of git://git.kernel.org/pub/scm/linux/kernel/git/net...	Linus Torvalds	3 days

Tag	Download	Author	Age
v6.0	net-next-6.0.tar.gz	Linus Torvalds	5 days
v6.0-rc7	net-next-6.0-rc7.tar.gz	Linus Torvalds	12 days
v6.0-rc6	net-next-6.0-rc6.tar.gz	Linus Torvalds	3 weeks
v6.0-rc5	net-next-6.0-rc5.tar.gz	Linus Torvalds	4 weeks
v6.0-rc4	net-next-6.0-rc4.tar.gz	Linus Torvalds	5 weeks
v6.0-rc3	net-next-6.0-rc3.tar.gz	Linus Torvalds	6 weeks
v6.0-rc2	net-next-6.0-rc2.tar.gz	Linus Torvalds	7 weeks
v6.0-rc1	net-next-6.0-rc1.tar.gz	Linus Torvalds	8 weeks
v5.19	net-next-5.19.tar.gz	Linus Torvalds	2 months
v5.19-rc8	net-next-5.19-rc8.tar.gz	Linus Torvalds	2 months
[...]			

Age	Commit message	Author	Files	Lines
3 days	Merge tag 'net-next-6.1' of git://git.kernel.org/pub/scm/linux/kernel/git/net... HEAD master main	Linus Torvalds	2136	-46141/+122889
3 days	Merge tag 'landlock-6.1-rc1' of git://git.kernel.org/pub/scm/linux/kernel/git...	Linus Torvalds	6	-54/+76
3 days	Merge tag 'audit-pr-20221003' of git://git.kernel.org/pub/scm/linux/kernel/gi...	Linus Torvalds	3	-13/+9
3 days	Merge tag 'x86_cleanups_for_v6.1_rc1' of git://git.kernel.org/pub/scm/linux/k...	Linus Torvalds	7	-12/+13
3 days	Merge tag 'x86_cache_for_v6.1_rc1' of git://git.kernel.org/pub/scm/linux/kern...	Linus Torvalds	8	-265/+523
3 days	Merge tag 'x86_microcode_for_v6.1_rc1' of git://git.kernel.org/pub/scm/linux/...	Linus Torvalds	6	-39/+119
3 days	Merge tag 'x86_paravirt_for_v6.1_rc1' of git://git.kernel.org/pub/scm/linux/k...	Linus Torvalds	2	-0/+2
3 days	Merge tag 'x86_misc_for_v6.1_rc1' of git://git.kernel.org/pub/scm/linux/kerne...	Linus Torvalds	2	-2/+2
3 days	Merge tag 'x86_asm_for_v6.1_rc1' of git://git.kernel.org/pub/scm/linux/kernel...	Linus Torvalds	1	-21/+33
3 days	Merge tag 'x86_core_for_v6.1_rc1' of git://git.kernel.org/pub/scm/linux/kerne...	Linus Torvalds	4	-8/+22
[...]				

문제 재현 및 원인 판단

```
$ cd samples/bpf
$ make -j
$ ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
```

커널의 테스트 스크립트
실행 도중 **에러 발생**

문제 재현 및 원인 판단

```
$ cd samples/bpf
$ make -j
$ ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
# vim tracex7_user.c
```

```
1 tracex7_user.c + X
int main(int argc, char **argv)
{
    // ... SNIP ...
    char filename[256];
    char command[256];
    int ret = 0;
    FILE *f;

    snprintf(filename, sizeof(filename), "%s_kern.o", argv[0]);
+--- 24 lines: obj = bpf_object__open_file(filename, NULL);-----

    snprintf(command, 256, "mount %s tmpmnt/", argv[1]);
    f = popen(command, "r");
    ret = pclose(f);

cleanup:
    bpf_link__destroy(link);
    bpf_object__close(obj);
    return ret ? 0 : 1;
}
NORMAL tracex7_user.c + 18% 8:0
```

커널의 테스트 스크립트
실행 도중 **에러 발생**

프로그램 실행 시 **인자를
확인하지 않아** 발생한 문제

버그 해결을 위해 소스코드 수정

```
$ git diff
```

```
diff --git a/samples/bpf/test_override_return.sh b/samples/bpf/test_override_return.sh
index e68b9ee6814b8b..35db26f736b9d9 100755
--- a/samples/bpf/test_override_return.sh
+++ b/samples/bpf/test_override_return.sh
@@ -1,5 +1,6 @@
#!/bin/bash
+rm -r tmpmnt
rm -f testfile.img
dd if=/dev/zero of=testfile.img bs=1M seek=1000 count=1
DEVICE=$(losetup --show -f testfile.img)
diff --git a/samples/bpf/tracex7_user.c b/samples/bpf/tracex7_user.c
index fdcd6580dd736a..8be7ce18d3ba05 100644
--- a/samples/bpf/tracex7_user.c
+++ b/samples/bpf/tracex7_user.c
@@ -14,6 +14,11 @@ int main(int argc, char **argv)
int ret = 0;
FILE *f;
+ if (!argv[1]) {
+ fprintf(stderr, "ERROR: Run with the btrfs device argument!\n");
+ return 0;
+ }
+
snprintf(filename, sizeof(filename), "%s_kern.o", argv[0]);
obj = bpf_object__open_file(filename, NULL);
if (libbpf_get_error(obj)) {
```

커밋 메시지 작성

소스 코드 수정 후, 커밋 메시지 작성 후 커밋 진행

```
$ git commit -s
```

커밋 메시지 작성

커밋 메시지 작성이 어렵다면?

```
$ git log -- <내가 수정한 파일 이름 filename>
```

커밋 메시지 작성

커밋 메시지 작성이 어렵다면?

```
$ git log -- samples/bpf/tracex7_user.c <해당 파일의 이전 커밋 확인>
```

```
commit 2bf3e2ef425bc2a164f10b554b7db6a8b4090ef4  
Author: Jakub Kicinski <kuba@kernel.org>  
Date: Mon May 14 22:35:02 2018 -0700
```

```
samples: bpf: include bpf/bpf.h instead of local libbpf.h
```

There are two files in the tree called libbpf.h which is becoming problematic. Most samples don't actually need the local libbpf.h they simply include it to get to bpf/bpf.h. Include bpf/bpf.h directly instead.

```
Signed-off-by: Jakub Kicinski <jakub.kicinski@netronome.com>  
Acked-by: Jesper Dangaard Brouer <brouer@redhat.com>  
Signed-off-by: Alexei Starovoitov <ast@kernel.org>
```

```
commit 965de87e54b803223bff703ea6b2a76c056695ae  
Author: Josef Bacik <jbacik@fb.com>  
Date: Mon Dec 11 11:36:49 2017 -0500
```

```
samples/bpf: add a test for bpf_override_return
```


커밋 메시지 작성

```
samples: bpf: Fix tracex7 error raised on the missing argument
```

The current behavior of 'tracex7' doesn't consist with other bpf samples `tracex{1..6}`. Other samples do not require any argument to run with, but `tracex7` should be run with `btrfs` device argument. (it should be executed with `test_override_return.sh`)

Currently, `tracex7` doesn't have any description about how to run this program and raises an unexpected error. And this result might be confusing since users might not have a hunch about how to run this program.

```
// Current behavior
# ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
// Fixed behavior
# ./tracex7 ERROR: Run with the btrfs device argument!
```

In order to fix this error, this commit adds logic to report a message and exit when running this program with a missing argument.

Additionally in `test_override_return.sh`, there is a problem with multiple `directory(tmpmnt)` creation. So in this commit adds a line with removing the directory with every execution.

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```

커밋 메시지 작성

```
samples: bpf: Fix tracex7 error raised on the missing argument
```

The current behavior of 'tracex7' doesn't consist with other bpf samples tracex{1..6}. Other samples do not require any argument to run with, but tracex7 should be run with btrfs device argument. (it should be executed with test_override_return.sh)

Currently, tracex7 doesn't have any description about how to run this program and raises an unexpected error. And this result might be confusing since users might not have a hunch about how to run this program.

```
// Current behavior
# ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
// Fixed behavior
# ./tracex7 ERROR: Run with the btrfs device argument!
```

In order to fix this error, this commit adds logic to report a message and exit when running this program with a missing argument.

Additionally in test_override_return.sh, there is a problem with multiple directory(tmpmnt) creation. So in this commit adds a line with removing the directory with every execution.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>

커밋 헤더는
가능하다면 50자 이내

커밋 메시지 작성

```
samples: bpf: Fix tracex7 error raised on the missing argument
```

The current behavior of 'tracex7' doesn't consist with other bpf samples tracex{1..6}. Other samples do not require any argument to run with, but tracex7 should be run with btrfs device argument. (it should be executed with test_override_return.sh)

Currently, tracex7 doesn't have any description about how to run this program and raises an unexpected error. And this result might be confusing since users might not have a hunch about how to run this program.

```
// Current behavior
# ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
// Fixed behavior
# ./tracex7 ERROR: Run with the btrfs device argument!
```

In order to fix this error, this commit adds logic to report a message and exit when running this program with a missing argument.

Additionally in test_override_return.sh, there is a problem with multiple directory(tmpmnt) creation. So in this commit adds a line with removing the directory with every execution.

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```

**커밋 설명은 상세하게
한 줄에 72자 넘지 않도록**

커밋 메시지 작성

```
samples: bpf: Fix tracex7 error raised on the missing argument
```

The current behavior of 'tracex7' doesn't consist with other bpf samples `tracex{1..6}`. Other samples do not require any argument to run with, but `tracex7` should be run with `btrfs` device argument. (it should be executed with `test_override_return.sh`)

Currently, `tracex7` doesn't have any description about how to run this program and raises an unexpected error. And this result might be confusing since users might not have a hunch about how to run this program.

```
// Current behavior
# ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
// Fixed behavior
# ./tracex7 ERROR: Run with the btrfs device argument!
```

In order to fix this error, this commit adds logic to report a message and exit when running this program with a missing argument.

Additionally in `test_override_return.sh`, there is a problem with multiple directory(`tmpmnt`) creation. So in this commit adds a line with removing the directory with every execution.

```
Signed-off-by: Juhee Kang <claudiajkang@gmail.com>
```

**필요하다면, 발생한 문제의
상세한 설명을 위해 재현방법 서술**

커밋 메시지 작성

```
samples: bpf: Fix tracex7 error raised on the missing argument
```

The current behavior of 'tracex7' doesn't consist with other bpf samples tracex{1..6}. Other samples do not require any argument to run with, but tracex7 should be run with btrfs device argument. (it should be executed with test_override_return.sh)

Currently, tracex7 doesn't have any description about how to run this program and raises an unexpected error. And this result might be confusing since users might not have a hunch about how to run this program.

```
// Current behavior
# ./tracex7
sh: 1: Syntax error: word unexpected (expecting ")")
// Fixed behavior
# ./tracex7 ERROR: Run with the btrfs device argument!
```

In order to fix this error, this commit adds logic to report a message and exit when running this program with a missing argument.

Additionally in test_override_return.sh, there is a problem with multiple directory(tmpmnt) creation. So in this commit adds a line with removing the directory with every execution.

Signed-off-by: Juhee Kang <claudiajkang@gmail.com>

Signed-off-by 태그 작성

커널 개발 환경 구성 - Patch 메일 전송을 위한 환경 구축

```
# apt-get install git-email  
  
$ git config --global sendemail.smtpencryption tls  
$ git config --global sendemail.smtpserver <mailserver>  
$ git config --global sendemail.smtpuser <email>  
$ git config --global sendemail.smtpserverport 587  
$ git config --global sendemail.smtppass <password>
```

git-email 패키지를 통해
작성된 patch를 쉽게 메일로 전송 가능

git config을 이용해 **smtp 계정 설정**
ex) gmail smtp 설정

패치 메일 작성 준비 및 컨벤션 테스트

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"
```

패치 메일 작성 준비 및 컨벤션 테스트

〈커밋 개수〉

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"
```


패치 메일 작성 준비 및 컨벤션 테스트

〈커밋 개수〉

〈커널 tree에 따른 메일 prefix 추가〉

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"
```

패치 메일 작성 준비 및 컨벤션 테스트

〈커밋 개수〉

〈커널 tree에 따른 메일 prefix 추가〉

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"  
0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch
```

패치 메일 작성 준비 및 컨벤션 테스트

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"  
0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch  
  
$ apt-get install codespell python-ply python-git  
$ ./scripts/checkpatch.pl --strict --codespell 0001-*.patch
```

패치 메일 작성 준비 및 컨벤션 테스트

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"  
0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch  
  
$ apt-get install codespell python-ply python-git  
$ ./scripts/checkpatch.pl --strict --codespell 0001-*.patch    <인자로 패치 파일 전달>
```

패치 메일 작성 준비 및 컨벤션 테스트

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"  
0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch  
  
$ apt-get install codespell python-ply python-git  
$ ./scripts/checkpatch.pl --strict --codespell 0001-*.patch <인자로 패치 파일 전달>  
total: 0 errors, 0 warnings, 0 checks, 17 lines checked  
  
0001-*.patch has no obvious style problems and is ready for submission.
```

패치 메일 작성 준비 및 컨벤션 테스트

```
$ git format-patch -1 --subject-prefix="PATCH, bpf-next"  
0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch  
  
$ apt-get install codespell python-ply python-git  
$ ./scripts/checkpatch.pl --strict --codespell 0001-*.patch <인자로 패치 파일 전달>  
total: 0 errors, 0 warnings, 0 checks, 17 lines checked  
  
0001-*.patch has no obvious style problems and is ready for submission.  
    <컨벤션 문제 없음>
```

패치 메일 수신 대상 찾기 및 전송

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats 0001-*.patch
```

메인테이너와
메일링 리스트 찾기

패치 메일 수신 대상 찾기 및 전송

〈인자로 패치 파일 전달〉

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats 0001-*.patch
```

메인테이너와
메일링 리스트 찾기

패치 메일 수신 대상 찾기 및 전송

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats 0001-*.patch  
Alexei Starovoitov <ast@kernel.org>, Daniel Borkmann <daniel@iogearbox.net>,  
Andrii Nakryiko <andrii@kernel.org>, netdev@vger.kernel.org, bpf@vger.kernel.org
```

메인테이너와
메일링 리스트 찾기

패치 메일 수신 대상 찾기 및 전송

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats 0001-*.patch  
Alexei Starovoitov <ast@kernel.org>, Daniel Borkmann <daniel@iogearbox.net>,  
Andrii Nakryiko <andrii@kernel.org>, netdev@vger.kernel.org, bpf@vger.kernel.org
```

```
$ git send-email --to="maintainer mail" --to="mailing list mail" 0001-*.patch  
Send this email? ([y]es|[n]o|[q]uit|[a]ll): Y
```

```
[...]
```

```
Result: OK
```

메일 전송하기

메일 전송되었는지 확인

<input type="checkbox"/>	★	▶	me	Inbox	[PATCH net-next 2/2] samples: pktgen: add missing IPv6 option to pktgen scripts - Currently, "sample04" and "sample05" are not working properly when running with an IPv...	8/13/21
<input type="checkbox"/>	★	▶	me	Inbox	[PATCH net-next 1/2] samples: pktgen: pass the environment variable of normal user to sudo - All pktgen samples can use the environment variable instead of option param...	8/13/21
<input type="checkbox"/>	★	▶	me	Inbox	[PATCH net-next 0/2] samples: pktgen: enhance the usability of pktgen samples - This patchset improves the usability of pktgen samples by adding an option for propagati...	8/13/21
<input type="checkbox"/>	☆	▶	Majordomo 2	Inbox	Confirmation for subscribe netdev - -- Someone (possibly you) has requested that your email address be added to or deleted from the mailing list "netdev@vger.kernel.org"....	8/5/21
<input type="checkbox"/>	★	▶	me .. patchwork-b. 6	Inbox	[bpf-next 1/2] samples: bpf: Fix tracex7 error raised on the missing argument - Hello: This series was applied to bpf/bpf-next.git (refs/heads/master): On Tue, 27 Jul 2021 04:1...	7/28/21
<input type="checkbox"/>	☆	▶	me	Inbox	[bpf-next v2 2/2] samples: bpf: Add the omitted xdp samples to .gitignore - There are recently added xdp samples (xdp_redirect_map_multi and xdpsock_ctrl_proc) which ar...	7/27/21
<input type="checkbox"/>	★	▶	me, Yonghong 2	Inbox	[bpf-next 2/2] samples: bpf: Add the omitted xdp samples to .gitignore - On 7/24/21 8:21 AM, Juhee Kang wrote: > There are recently added xdp samples (xdp_redirect_map_...	7/26/21

Gmail을 통해서도 확인 가능

메일 전송되었는지 확인

Show patches with: Submitter = Juhee Kang | 29 patches

Patch	Series	A/R/T	S/W/F	▲ Date	Submitter	Delegate	State
[net-next,3/3] pktgen: document the latest pktgen usage options	samples: pktgen: enhance the ability to print the execution results of samples	---	14 1 -	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/3] samples: pktgen: add trap SIGINT for printing execution result	samples: pktgen: enhance the ability to print the execution results of samples	---	14 1 -	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,1/3] samples: pktgen: fix to print when terminated normally	samples: pktgen: enhance the ability to print the execution results of samples	---	14 1 -	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/2] samples: pktgen: add missing IPv6 option to pktgen scripts	samples: pktgen: enhance the usability of pktgen samples	---	13 - 2	2021-08-12	Juhee Kang	netdev	Accepted
[net-next,1/2] samples: pktgen: pass the environment variable of normal user to sudo	samples: pktgen: enhance the usability of pktgen samples	---	14 1 -	2021-08-12	Juhee Kang	netdev	Accepted
[bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	---	14 1 -	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	1 --	14 1 -	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	1 --	14 1 -	2021-07-24	Juhee Kang	bpf	Superseded
[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	1 --	14 1 -	2021-07-24	Juhee Kang	bpf	Superseded

<https://patchwork.kernel.org/project/netdevbpf/list>

Patchwork을 통해서도 확인

리뷰 확인

```
From: Andrii Nakryiko <andrii.nakryiko@gmail.com>  
To: Juhee Kang <claudiajkang@gmail.com>  
Cc: Daniel Borkmann <daniel@iogearbox.net>, Alexei Starovoitov <ast@kernel.org>,  
    Andrii Nakryiko <andrii@kernel.org>, Networking <netdev@vger.kernel.org>, bpf <bpf@vger.kernel.org>  
Subject: Re: [bpf-next 1/2] samples: bpf: Fix tracex7 error raised on the missing argument
```

```
On Sat, Jul 24, 2021 at 8:21 AM Juhee Kang <claudiajkang@gmail.com> wrote:
```

```
[...]  
> diff --git a/samples/bpf/test_override_return.sh b/samples/bpf/test_override_return.sh  
> index e68b9ee6814b..6480b55502c7 100755  
> --- a/samples/bpf/test_override_return.sh  
> +++ b/samples/bpf/test_override_return.sh  
> @@ -1,5 +1,6 @@  
> #!/bin/bash  
>  
> +rm -rf tmpmnt
```

```
Do we need -rf or -r would do?
```

```
[...]  
> --  
> 2.27.0
```

리뷰 응답

```
From: Juhee Kang <claudiajkang@gmail.com>
To: Andrii Nakryiko <andrii.nakryiko@gmail.com>
Cc: Daniel Borkmann <daniel@iogearbox.net>, Alexei Starovoitov <ast@kernel.org>,
    Andrii Nakryiko <andrii@kernel.org>, Networking <netdev@vger.kernel.org>, bpf <bpf@vger.kernel.org>
Subject: Re: [bpf-next 1/2] samples: bpf: Fix tracex7 error raised on the missing argument
```

```
On Tue, Jul 27, 2021 at 5:08 AM Andrii Nakryiko <andrii.nakryiko@gmail.com> wrote:
```

```
> On Sat, Jul 24, 2021 at 8:21 AM Juhee Kang <claudiajkang@gmail.com> wrote:
```

```
[...]
> > diff --git a/samples/bpf/test_override_return.sh b/samples/bpf/test_override_return.sh
> > index e68b9ee6814b..6480b55502c7 100755
> > --- a/samples/bpf/test_override_return.sh
> > +++ b/samples/bpf/test_override_return.sh
> > @@ -1,5 +1,6 @@
> > #!/bin/bash
> >
> > +rm -rf tmpmnt
```

```
> Do we need -rf or -r would do?
```

```
It works properly using only -r.
```

```
Thanks for pointing it out! I will stick to this method.
```

```
I will send the next version as soon as possible.
```

```
[...]
```

```
> > 2.27.0
```

코드 수정 후 패치 메일 작성 준비

〈두번째 버전의 패치임을 명시〉

```
$ git format-patch -1 -v2 --subject-prefix="PATCH, bpf-next"
```

코드 수정 후 패치 메일 작성 준비

```
$ git format-patch -1 -v2 --subject-prefix="PATCH, bpf-next"  
v2-0001-samples-bpf-Fix-tracex7-error-raised-on-the-missing-.patch  
  
$ ./scripts/checkpatch.pl --strict --codespell v2-0001-*.patch  
total: 0 errors, 0 warnings, 0 checks, 17 lines checked  
  
v2-0001-*.patch has no obvious style problems and is ready for submission.
```


패치 메일 수신 대상 찾기 및 전송

```
$ ./scripts/get_maintainer.pl --no-multiline --no-rolestats v2-0001-*.patch  
Alexei Starovoitov <ast@kernel.org>, Daniel Borkmann <daniel@iogearbox.net>,  
Andrii Nakryiko <andrii@kernel.org>, netdev@vger.kernel.org, bpf@vger.kernel.org
```

```
$ git send-email --to="maintainer mail" --to="mailing list mail" v2-0001-*.patch  
Send this email? ([y]es|[n]o|[q]uit|[a]ll): Y
```

```
[...]
```

```
Result: OK
```

메일 전송하기

패치 accept

```
From: patchwork-bot+netdevbpf <patchwork-bot+netdevbpf@kernel.org>  
To: Juhee Kang <claudiajkang@gmail.com>  
Subject: Re: [bpf-next v2 1/2] samples: bpf: Fix tracex7 error raised on the missing argument
```

Hello:

This series was applied to [bpf/bpf-next.git \(refs/heads/master\)](#):

On Tue, 27 Jul 2021 04:10:55 +0000 you wrote:

```
> The current behavior of 'tracex7' doesn't consist with other bpf samples  
> tracex{1..6}. Other samples do not require any argument to run with, but  
> tracex7 should be run with btrfs device argument. (it should be executed  
> with test_override_return.sh)  
> [...]
```

Here is the summary with links:

- [bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument
<https://git.kernel.org/bpf/bpf-next/c/7d07006f0592>
- [bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore
<https://git.kernel.org/bpf/bpf-next/c/05e9b4f60d31>

You are awesome, thank you!

--

Deet-doot-dot, I am a bot.
<https://korg.docs.kernel.org/patchwork/pwbot.html>

패치 accept



patchwork-bot+netdevbpf@kernel.org

to me, daniel, ast, andrii, yhs, netdev, bpf ▾

Hello:

This series was applied to bpf/bpf-next.git (refs/heads/master):

On Tue, 27 Jul 2021 04:10:55 +0000 you wrote:

```
> The current behavior of 'tracex7' doesn't consist with other bpf samples
> tracex{1..6}. Other samples do not require any argument to run with, but
> tracex7 should be run with btrfs device argument. (it should be executed
> with test_override_return.sh)
>
> Currently, tracex7 doesn't have any description about how to run this
> program and raises an unexpected error. And this result might be
> confusing since users might not have a hunch about how to run this
> program.
>
> [...]
```

Here is the summary with links:

- [bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument
<https://git.kernel.org/bpf/bpf-next/c/7d07006f0592>
- [bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore
<https://git.kernel.org/bpf/bpf-next/c/05e9b4f60d31>

You are awesome, thank you!

--

Deet-doot-dot, I am a bot.

<https://korg.docs.kernel.org/patchwork/pwbot.html>

Gmail을 통해서도 확인 가능

패치 accept

Patchwork Netdev + BPF Patches Bundles About this Login Register Mail settings

Show patches with: Submitter = Juhee Kang | 29 patches

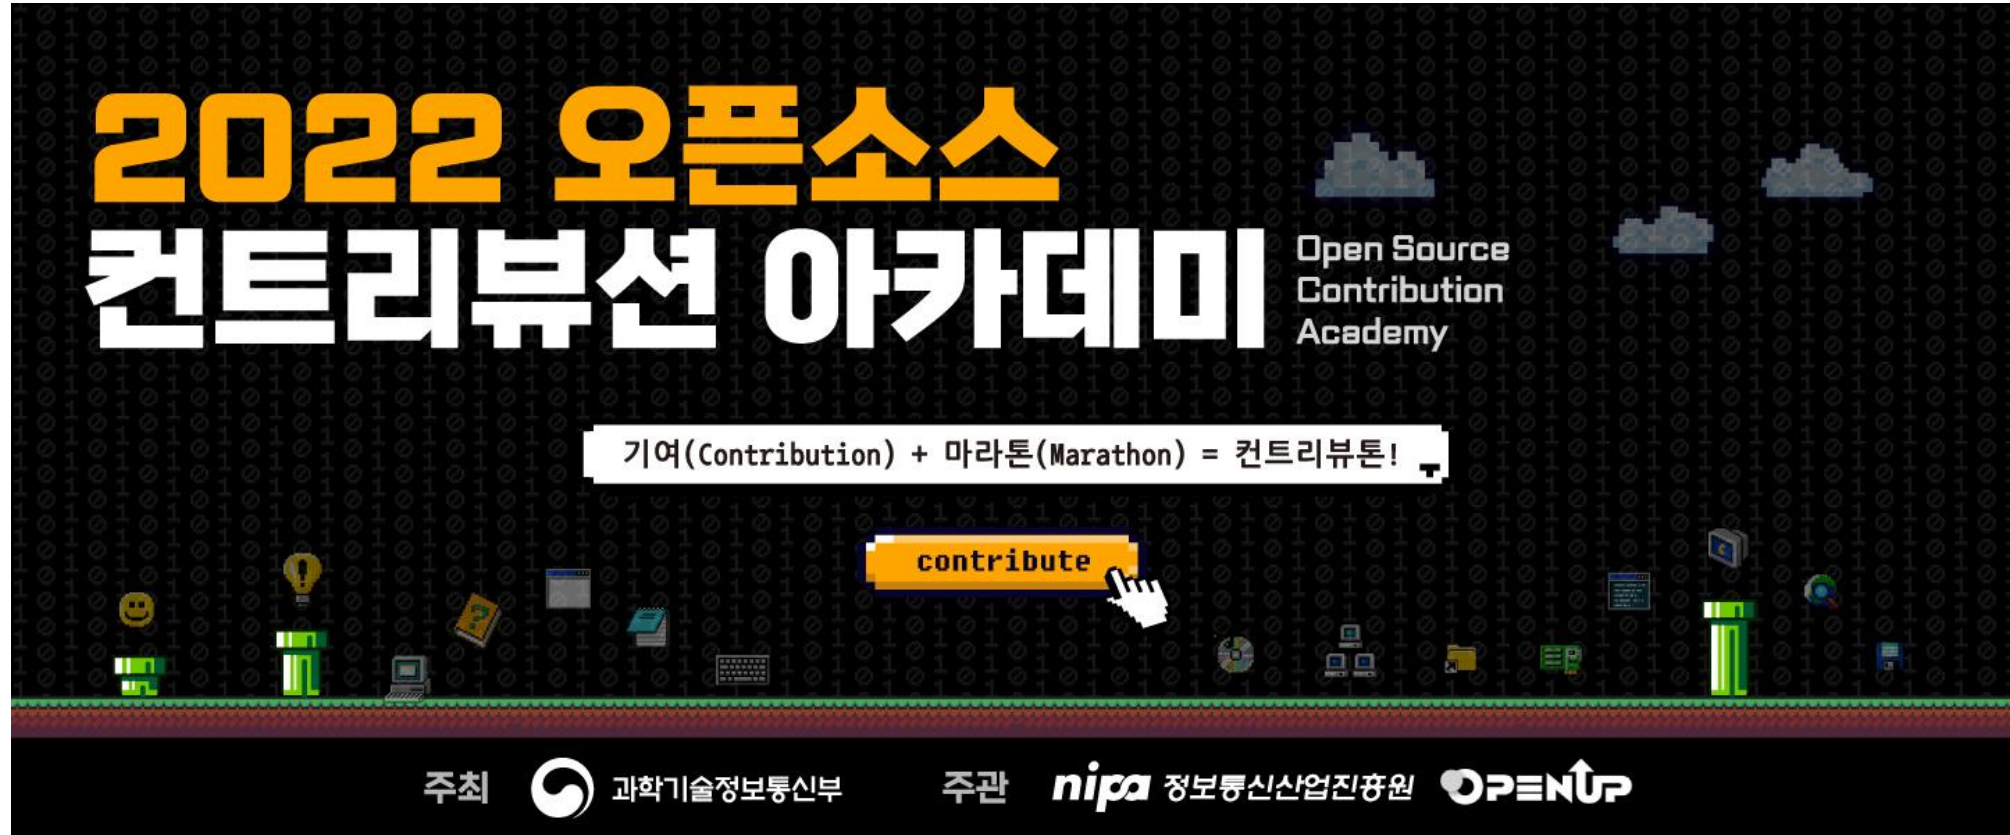
Patch	Date	Submitter	Delegate	State
[net-next,3/3] pktgen: document the latest pktgen usage options	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/3] samples: pktgen: add trap SIGINT for printing execution result	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,1/3] samples: pktgen: fix to print when terminated normally	2021-08-25	Juhee Kang	netdev	Accepted
[net-next,2/2] samples: pktgen: add missing IPv6 option to pktgen scripts	2021-08-12	Juhee Kang	netdev	Accepted
[net-next,1/2] samples: pktgen: pass the environment variable of normal user to su	2021-08-12	Juhee Kang	netdev	Accepted
[bpf-next,v2,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,v2,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-27	Juhee Kang	bpf	Accepted
[bpf-next,2/2] samples: bpf: Add the omitted xdp samples to .gitignore	2021-07-24	Juhee Kang	bpf	Superseded
[bpf-next,1/2] samples: bpf: Fix tracex7 error raised on the missing argument	2021-07-24	Juhee Kang	bpf	Superseded

<https://patchwork.kernel.org/project/netdevbpf/list>

Patchwork을 통해서도 확인


여기까지가 간단한 리눅스 커널 커밋 과정이었습니다.

커널 패치를 진행해보고 싶으시다면?



The banner features a dark background with a repeating pattern of small circles. At the top left, the text "2022 오픈소스" is written in large, bold, orange Korean characters. Below it, "컨트리뷰션 아카데미" is written in large, bold, white Korean characters. To the right of this, the English text "Open Source Contribution Academy" is displayed in a smaller, white font. In the center, a white box contains the Korean text "기여(Contribution) + 마라톤(Marathon) = 컨트리뷰톤!". Below this, a yellow button with the word "contribute" in black text is shown with a white hand cursor pointing at it. The bottom of the banner features a row of various icons representing different aspects of technology and development, such as a smiley face, a lightbulb, a laptop, a document, a keyboard, a globe, and a server. At the very bottom, there are logos for the organizing institutions: "주최 과학기술정보통신부" (Organized by the Ministry of Science and ICT), "주관 nipa 정보통신산업진흥원" (Organized by NIPA, Korea Communications Commission), and the "OPENUP" logo.

커널 관련 컨트리뷰션 프로그램 참여



커널 패치를 위한 장벽 중 하나라도 극복이 되셨길!

Q & A