

Considerations for Benchmarking Network Performance in Containerized Infrastructure

draft-ietf-bmwg-containerized-infra-06

Minh-Ngoc Tran (Soongsil University), Sridhar Rao (The Linux Foundation),
Jangwon Lee, Younghan Kim (Soongsil University)

Scope Re-introduction

- Previous **NFV benchmarking** related RFCs
 - RFC 8172: Considerations for Benchmarking Virtual Network Functions and Their Infrastructure
 - RFC 8204: Benchmarking Virtual Switches in the Open Platform for NFV (OPNFV)

• The primary scope of this document is to fill in **the gaps** of these works when applying to **containerized NFV** infrastructure.

- **The consideration gaps are:**

- Different **network models/topologies configured by Container Network Interfaces** (including the extended Berkeley Packet Filter model which was not mentioned in previous documents)
- **Resources configuration for containers.**

4.1. Networking Models	5	4.2. Resources Configuration	15
4.1.1. Kernel-space non-Acceleration Model	6	4.2.1. CPU Isolation / NUMA Affinity	15
4.1.2. User-space Acceleration Model	7	4.2.2. Pod Hugepages	16
4.1.3. eBPF Acceleration Model	8	4.2.3. Pod CPU Cores and Memory Allocation	16
4.1.4. Smart-NIC Acceleration Model	13	4.2.4. Service Function Chaining	17
4.1.5. Model Combination	14		

Updates (1)

Clearer categorization of eBPF Acceleration Model

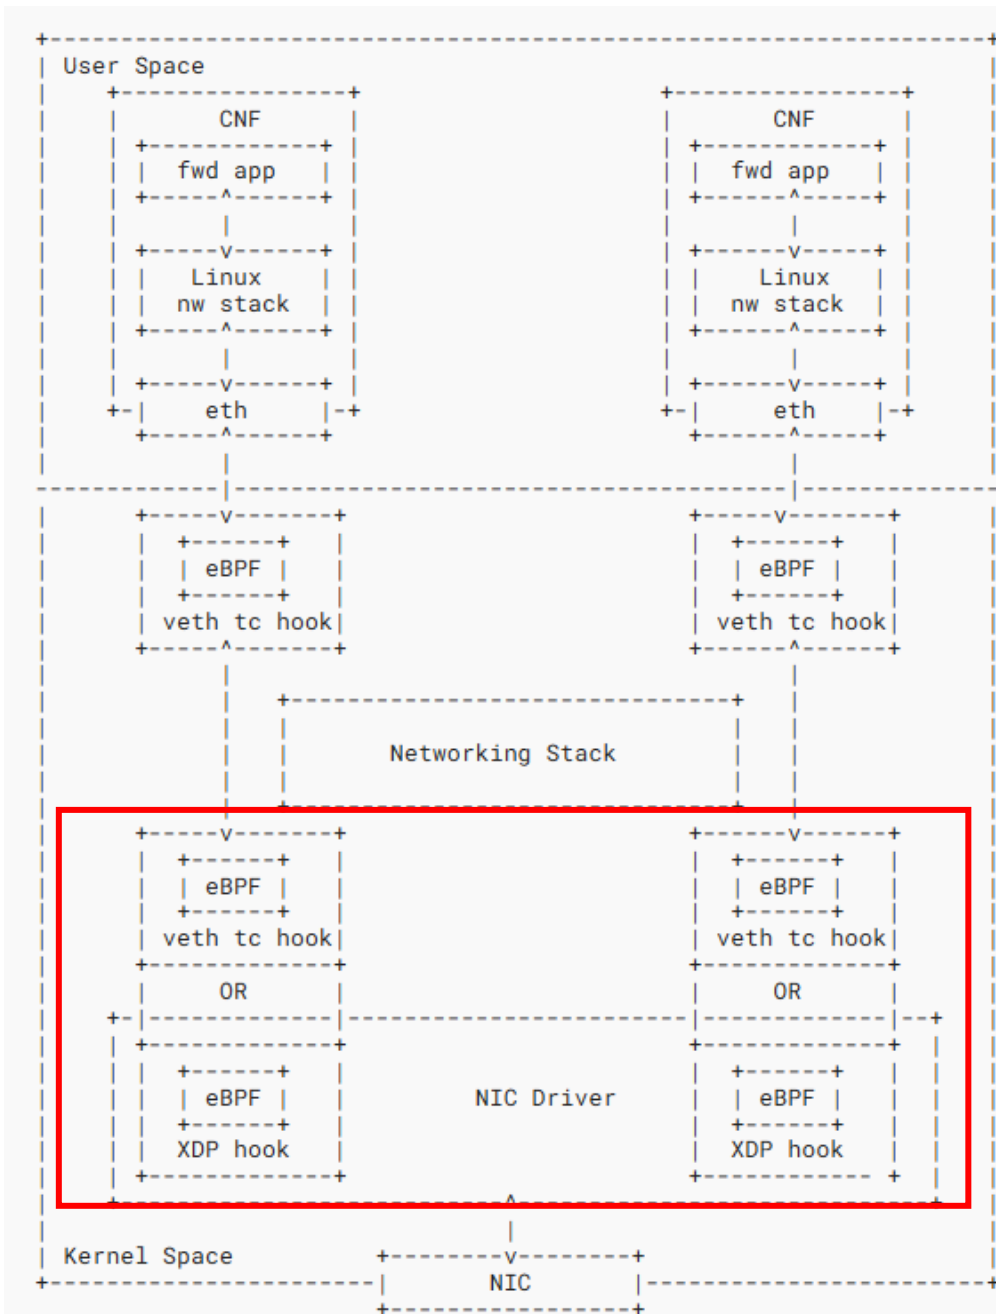
Previous version

- Categorizes based on packet acceleration path **using AF_XDP or not, and the enabling technology for AF_XDP**
 - **Non-AF_XDP**
 - Packet go through traffic control/XDP hook but **not go through AF_XDP port**
 - **Using AF_XDP supported CNI**
 - Packet go through XDP hook **to AF_XDP port controlled by the AF_XDP supported CNI**
 - **Using AF_XDP supported vSwitch**
 - Packet go through XDP hook **to AF_XDP port controlled by vSwitch**, vSwitch-Pod is controlled by CNI

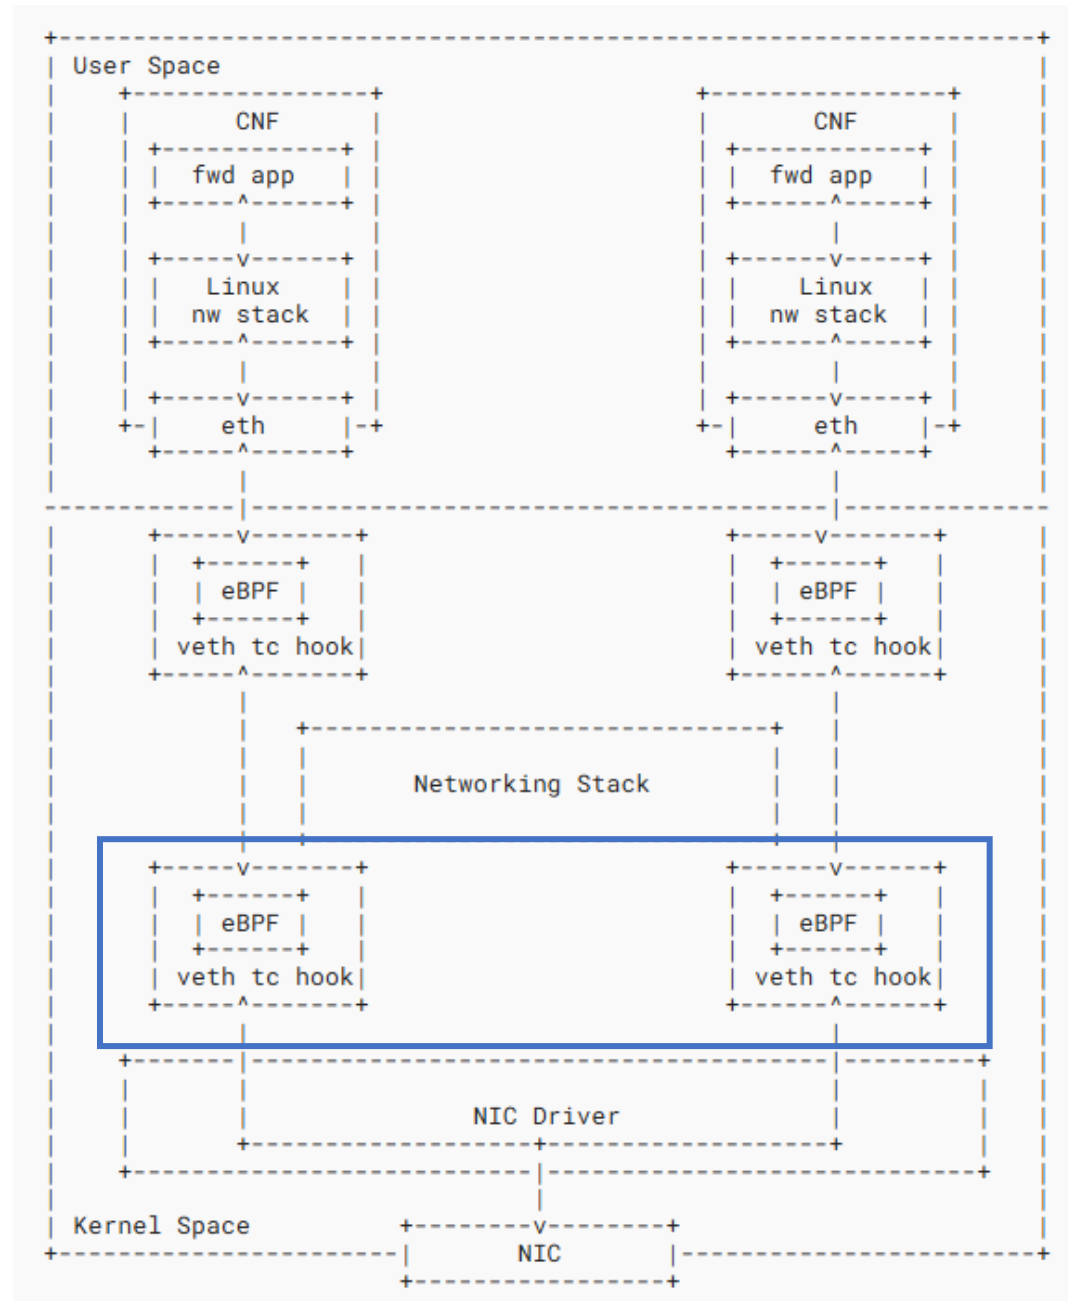
➔ Latest version

- Categorizes based on packet acceleration path **going through which kind of eBPF hook**
 - **Forwarding via traffic control hook**
 - Packet **go through eBPF traffic control hook**
 - **Forwarding via XDP hook**
 - Packet **go through eBPF XDP hook to AF_XDP port**
 - Different methods to reach pod based on the implemented CNI

Non-AF_XDP



-> Forward via traffic control hook



Examples

- Calico
- Default Cilium

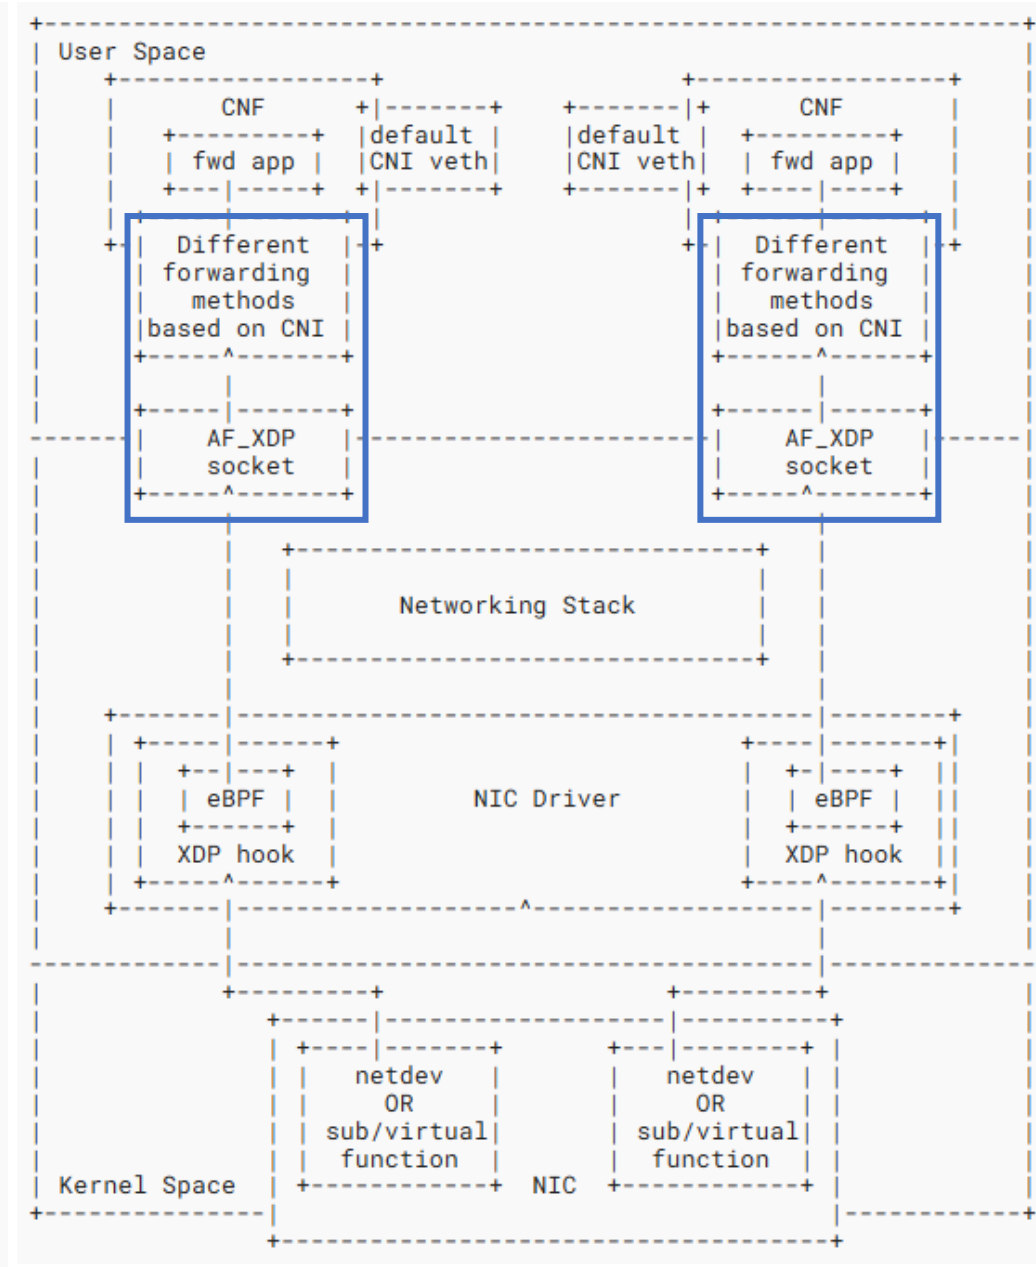
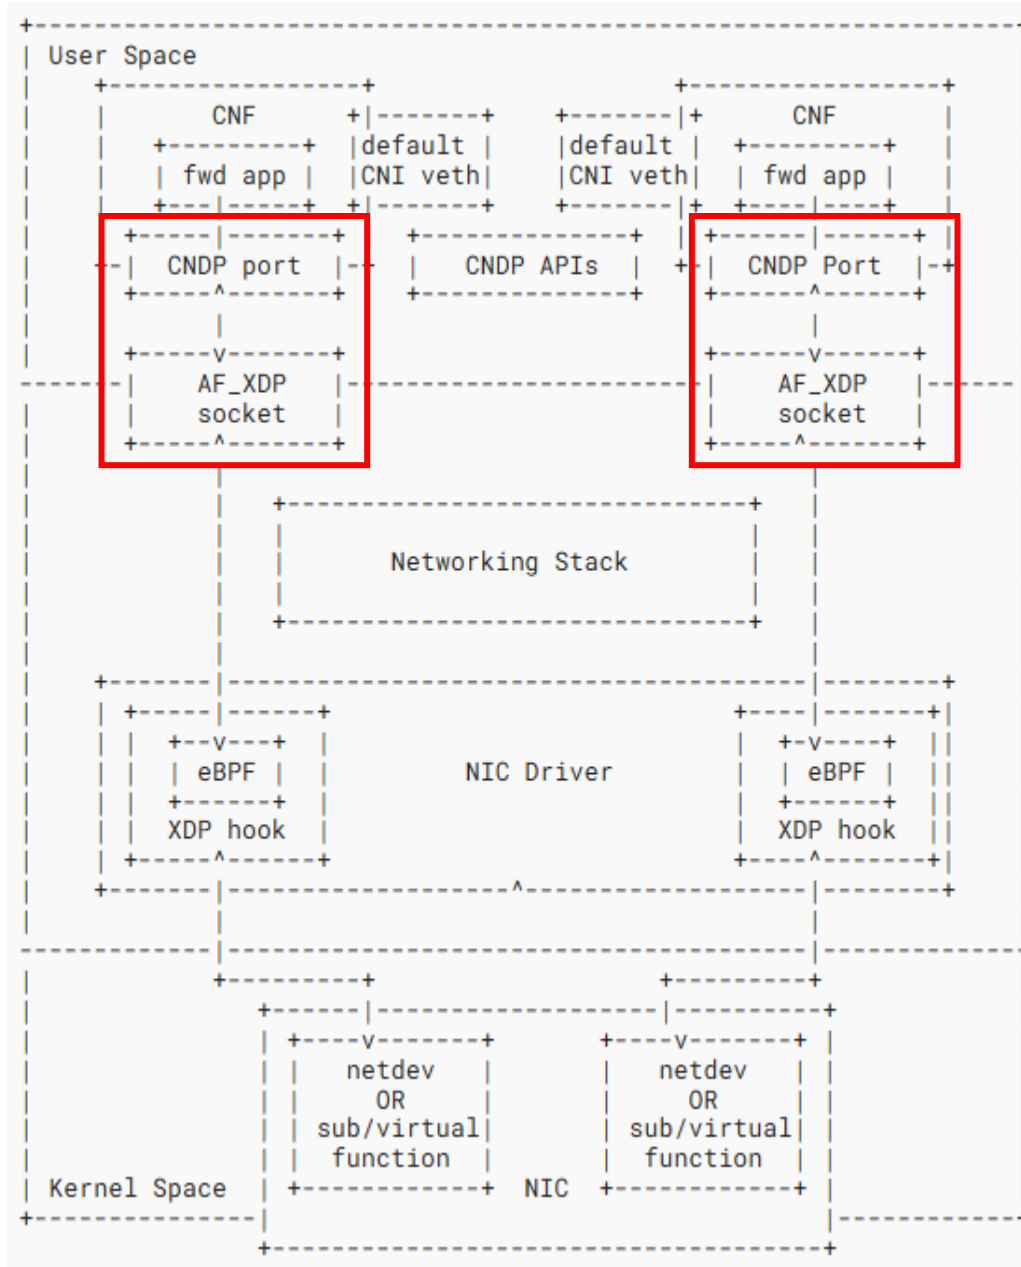
XDP-enabled Cilium: packets still go through AF_XDP port, which means the previous figure is incorrect

Using AF_XDP supported CNI

-> Forward via XDP hook

Examples

- AF_XDP K8s CNI Plugin
- AF_XDP supported vSwitch (OVSDPDK/VPP)
- XDP-enabled Cilium

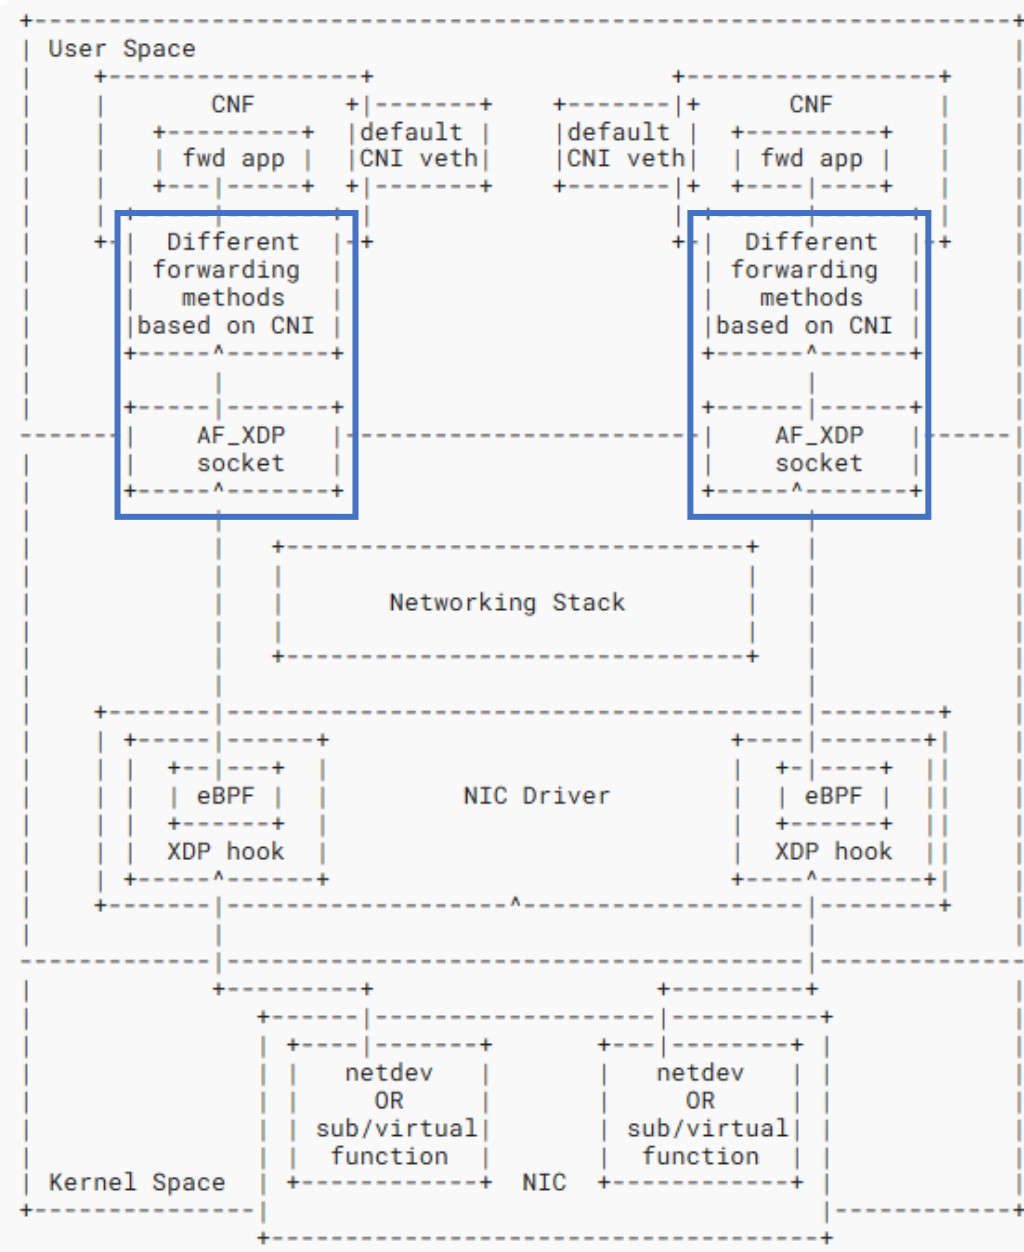
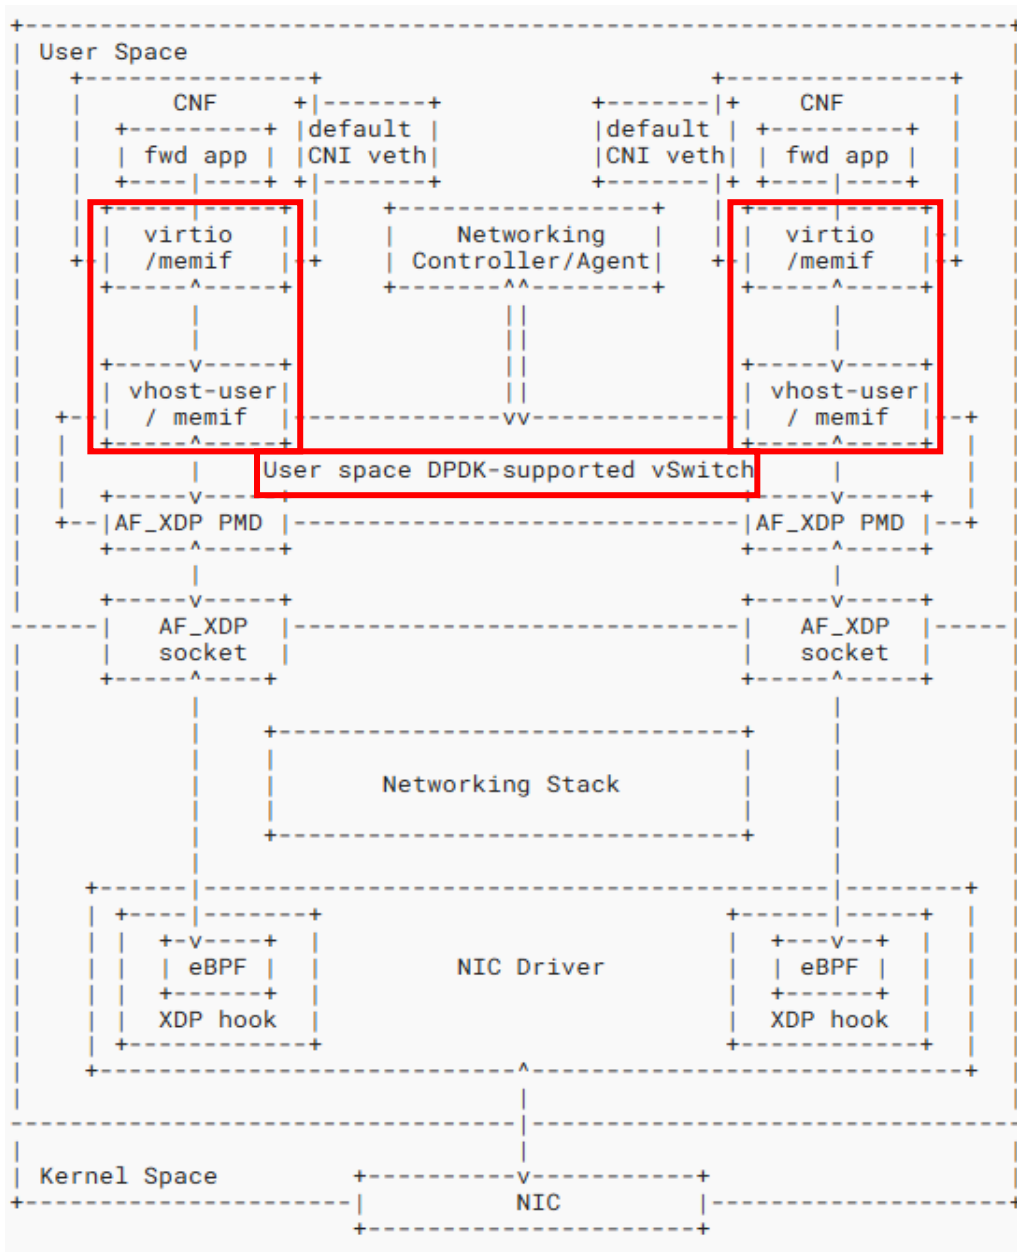


Using AF_XDP supported vSwitch

-> Forward via XDP hook

Examples

- AF_XDP K8s CNI Plugin
- AF_XDP supported vSwitch (OVSDPDK/VPP)
- XDP-enabled Cilium



Updates (1)

Clearer categorization of eBPF Acceleration Model

- Clearer comparison for reader to understand
- Correctly categorizes Cilium
 - default version via traffic control hook
 - XDP-enabled version via XDP hook
- **These changes are agreed by related RFC 8204 author - Maryam**
- **Forwarding via traffic control hook**
 - Packet **go through eBPF traffic control hook**
- **Forwarding via XDP hook**
 - Packet **go through eBPF XDP hook to AF_XDP port**
 - Different methods to reach pod based on the implemented CNI

Updates (2)

- Major review and enhance English grammar of the document
- Update references with their latest information

Summary

- The document is now technical and format Ready
- We would like to request Last Call.