

AF_XDP

PR #3314 .

Next (from 25Jan RM call):

1. How applications adopt to AF_XDP? (Now in text below)
2. How could it work in virtualized environments? (Now in text below)
3. How could it work on the same node with another app using DPDK/SR-IOV?

Text WIP before creating PR:

Chapter 3 (before Energy Efficiency)

AF_XDP

[Address Family For XDP \(AF_XDP\)](#) is optimized for high performance packet processing and introduced in Linux kernel v4.18. This new socket type leverages the [eXpress Data Path \(XDP\)](#) in-kernel fast-path to transfer traffic frames from the NIC driver directly to userspace without the need for full network stack. XDP is an [Extended Berkley Packet Filter \(eBPF\)](#) software program.

By using the XDP_REDIRECT action from that XDP program, ingress frames can be redirected to other XDP-enabled network devices. The fastest working mode of operation is Zero-Copy mode in enabled XDP drivers.

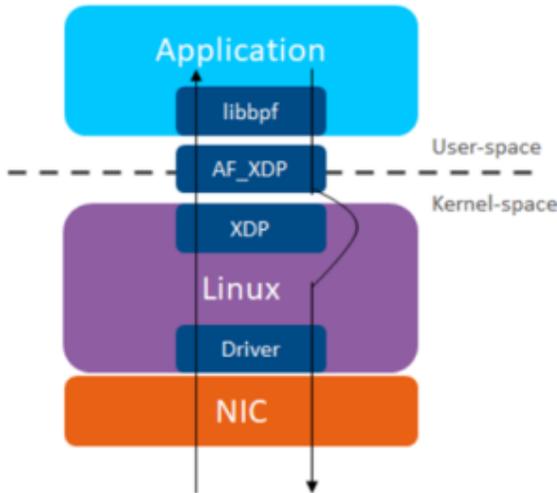


Figure X: AF_XDP architecture

Linux-native applications can open an AF_XDP socket to receive raw packets directly from the NIC, by using [libbpf](#) library functions to register a packet buffer area where packets will be located, and to create and bind the socket to a networking interface. DPDK-based applications can use [AF_XDP Poll Mode Driver](#). VPP-based applications can use [AF_XDP Device Driver](#).

In virtualized environments AF_XDP could be used as interface between guest Kernel and user space application, but still need SR-IOV or virtio to get traffic to the VM.

Chapter 5

Exposed Performance Optimisation Capabilities

Ref	Cloud Infrastructure Capability	Unit	Definition/Notes
e.cap.025	AF_XDP	Yes/No	Indicates if Cloud Infrastructure supports AF_XDP

Table 4-2: Exposed Performance Optimisation Capabilities of Cloud Infrastructure

Profiles Specifications & Capability Mapping

Ref	Capability	Basic	High Performance	Notes
e.cap.025	AF_XDP	No	Optional	These Capabilities require workload support for AF_XDP socket type.

Chapter 10 (last)

AF_XDP

Linux-native AF_XDP promises high enough packet processing performance and simplification compared to what SR-IOV and DPDK require for initial installation and later lifecycle management. Still, it will take time till AF_XDP-based solutions are financially invested and matured enough in both Virtualization Infrastructure and Network Functions.