

Alveo™ U25N SmartNIC: 25Gb Network and Security Acceleration

OVS with IPsec, Connection Tracking and Link Aggregation Acceleration

INTRODUCTION

Over the past few years, in server solutions, it is a well-established fact that network and storage functions are best offloaded onto SmartNICs. This architecture has been shown to be energy efficient and performance friendly. Energy efficiency comes from the fact that processing network and storage traffic at tens of gigabits per second consumes significant CPU resources which can be offloaded. In addition, SmartNIC offloading enables performance scalability to hundreds of gigabits per second of IO bandwidth.

As shown in Figure 3, when OVS + IPsec is running on the host CPU at large packet sizes, the maximum throughput achievable is limited to less than 4Gb/s. When OVS with IPsec are offloaded to the Alveo™ U25N SmartNIC with the same configuration, full line rate is supported for mid- and large packet sizes despite an extra tunneling header and IPsec header.

REAL-TIME ADAPTABLE HARDWARE ACCELERATION

Given the heterogeneity of workloads within network and storage across enterprise, edge, cloud and telco markets, one of the key attributes of the solution is to be adaptable. This implies flexibility in on-loading required functionality like Open vSwitch acceleration, IPsec acceleration, link aggregation, connection tracking etc. onto a particular FPGA during runtime.

This adaptability also enables removing unnecessary functions not needed for a given deployment. This in turn saves costs, power, testing and debug costs for the features of interest. Given this can be a particular FPGA instance specific choice, even within a larger deployment, various instances can enable various functions.

ALVEO U25N – DUAL PORT 25GB SMARTNIC

Given the heterogeneity of workloads across market segments, the Alveo U25N SmartNIC offers many functional IP choices included as part of its offering. In addition, customers can work with 3rd party developers to embellish other custom capabilities relevant to a deployment.

The Alveo U25N SmartNIC is based on the fusion of technologies, a powerful System on a Chip (SoC) that includes FPGA fabric and a multi-core Arm® processor, and an XtremeScale™ X2 Ethernet Controller. The FPGA enables hardware acceleration and offload to happen inline with maximum efficiency. The Arm cores are responsible for managing FPGA and control plane for the accelerations. The X2 Ethernet Controller provides a basic NIC function for two ports of 25 Gigabit Ethernet via software drivers.



KEY BENEFITS

- > Improving CPU efficiency with hardware-based offload accelerating networking and security functions
- > Delivering line-rate performance with turnkey accelerations including concurrent Open vSwitch (OVS), IPsec, and Firewall
- > Enabling flexibility to add custom plug-ins and workloads with application customization programmability

ALVEO™ U25N TURNKEY ACCELERATION

The Alveo U25N has three computational components, the X2 Ethernet Controller, the Arm® cores represented by the processing subsystem (PS) block, and the FPGA programmable logic (PL).

Most of the basic NIC processing tasks of the Alveo U25N are handled within the X2 controller, while the advanced functions are handled within the FPGA. Both blocks were designed to handle wire-rate traffic at 25Gb/s per port. The X2 controller processes standard network tasks and Single Root Input Output Virtualization (SR-IOV) requests to set up physical functions (PFs), virtual functions (VFs), and Receive Side Scaling (RSS).

The FPGA is home to several data plane pipeline stages. These include the OVS Engine which handles network virtualization tasks like VXLAN, L2GRE, IPsec, and Firewall.

In addition, OVS Connection tracking acceleration of stateful and stateless protocols TCP, UDP, ICMP etc. helps protect hypervisor from unwanted traffic and saves packet processing costs.

The Internet Protocol Security (IPsec) block on the FPGA offloads compute-intensive cryptographic functions, and it provides the open-source Internet Key Exchange (IKE) management solution - strongSwan. There is also a stateless firewall executing within the FPGA that can support rules passed down from the hosts' nftables entries.

The Arm cores are primarily used for control plane management of the FPGA's programmable logic. Applications running on the Arm cores update tables used by the OVS, IPsec, and firewall functions, while these programmable logic blocks are running. They also gather various statistics which are then used to improve the overall processing within the U25N SmartNIC platform.

With multiple physical and virtual ports, U25N offers bonds support with Link Aggregation Control protocol (LACP) capability that enables more than one port to be bonded together to form a logical port. This offers resiliency for link failures.

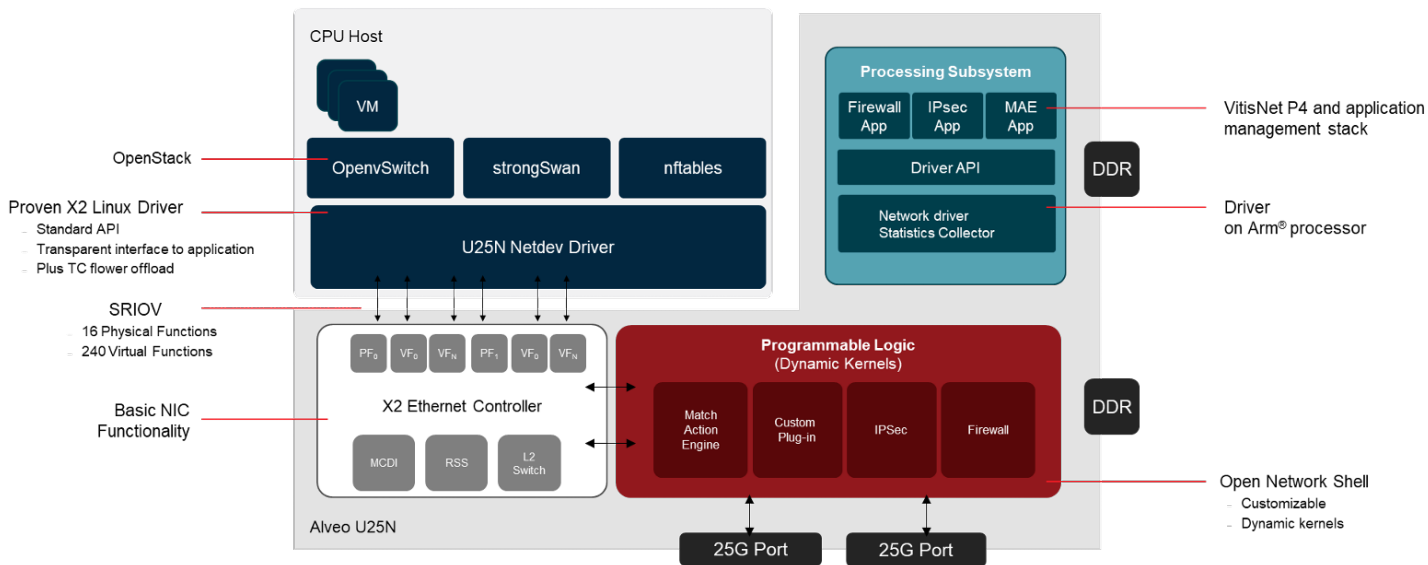


Figure 1: Alveo U25N Architectural Block Diagram

OVS ACCELERATION

The Alveo™ U25N supports full data plane offload and acceleration of virtual switching from the host CPU onto the SmartNIC. This offload enables a high-performance data path that seamlessly integrates with the control plane. All the packet processing operations such as packet classification, match/action/modification processing, and tunneling encapsulation/decapsulation are handled completely on the SmartNIC.

Below is a comparison demonstrating the benefits of offloading OVS onto the SmartNIC vs. running OVS in software on the host CPU. Figure 2 shows when OVS is run on the host CPU vs the Alveo U25N accelerator. There is a 4.4x difference in line-rate throughput and a two-third's reduction in CPU utilization. Furthermore, by offloading OVS to the SmartNIC, no CPU resources are consumed for OVS data plane processing.




Figure 2: OVS on Host CPU vs. OVS offloaded on the U25N SmartNIC

OVS + IPSEC ACCELERATION

In addition to OVS acceleration, data centers and enterprises also want to enable the secure transmission and receipt of network traffic between two sites. The Alveo U25N supports the concurrent operation of offloaded OVS and offloaded IPsec traffic without sacrificing performance.

IPsec tunnels are created between two servers. When IPsec is in transport mode, L2GRE is used to create tunnels. The strongSwan application runs in the user space. The strongSwan plugin is used to offload rules onto the Alveo U25N SmartNIC.

As shown in Figure 3, when OVS + IPsec is running on the host CPU at large packet sizes, the maximum throughput achievable was limited to less than 4Gb/s. When OVS + IPsec was offloaded to the Alveo U25N SmartNIC with the same configuration, full line rate is supported for mid and large packet sizes despite an extra tunneling header and IPsec header.

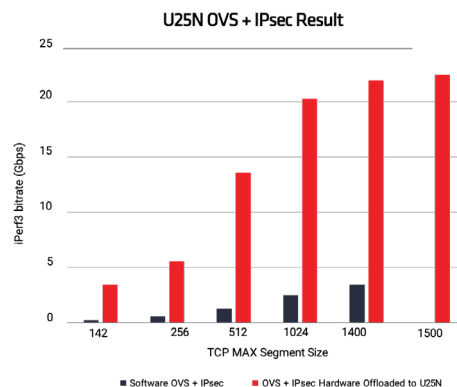


Figure 3: OVS + IPsec on Host CPU vs. OVS + IPsec offloaded on the U25N SmartNIC

IMPROVING INFRASTRUCTURE EFFICIENCY

The U25N SmartNIC enables improved CPU efficiency with hardware-based offload accelerating networking and security acceleration. This SmartNIC delivers line-rate performance with turnkey accelerations including concurrent Open vSwitch (OVS), IPsec, and Firewall.

TAKE THE NEXT STEP

- > Learn more about the Alveo U25N SmartNIC at www.xilinx.com/u25n
- > Learn more about AMD adaptive SmartNIC solutions at www.xilinx.com/network-acceleration
- > Learn more about the Alveo portfolio at www.xilinx.com/alveo
- > For additional product inquiries email dc_inquiries@amd.com

1: An average of 4.5x the performance [averaging 3.5x or 353% faster performance], using 67% fewer CPU cores using Alveo U25N SmartNIC with acceleration using driver v8.3, Open vSwitch v3.0.0 and DPDK v22.07 (512-byte packet length) than with same version of Open vSwitch and DPDK without acceleration. ALV-001 Testing conducted by AMD AECG DCCG Data Center Solution Architect Team as of 10/26/2022 on the Alveo U25N, on a test system comprising Model: Dell R740, Dual Sockets server, CPU: Intel® Xeon® Gold 6134 CPU @ 3.20GHz 8cores/16 HW threads, RAM: 96GB@2400 MT/s, HDD: 2.0TB, OS version: 5.15.0-46-generic #49~20.04.1-Ubuntu. PC manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers.

DISCLAIMERS

(The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale.

COPYRIGHT NOTICE

© Copyright 2023 Advanced Micro Devices, Inc. All rights reserved. Xilinx, the Xilinx logo, AMD, the AMD Arrow logo, Alveo, Artix, Kintex, Kria, Spartan, Versal, Vitis, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. AMBA, AMBA Designer, ARM, ARM1176JZ-S, CoreSight, Cortex, and PrimeCell are trademarks of ARM in the EU and other countries. PCIe, and PCI Express are trademarks of PCI-SIG and used under license. PID1847600

