TECHNICAL WHITE PAPER: **September 2024**

Accelerating Telecom Network Performance with DPU Offloading

6WIND virtual security gateway boosted by NVIDIA Data Processing Unit offloading

vmware® by Broadcom | NVIDIA | 6WIND

Table of contents

Introduction

Security gateways play a pivotal role in securing fixed and mobile communication networks by safeguarding the confidentiality of end-user payloads through data encryption and contributing to overall network integrity through the authentication of network elements.

Virtualizing security gateway functions should benefit communication service providers (CSPs) by lowering their total cost of ownership (TCO) and by enhancing deployment scalability, flexibility, and agility while delivering high performance and resiliency to support large mission-critical telco cloud deployments.

VMware Telco Cloud Platform

VMware Telco Cloud Platform™ is powered by field-proven compute, a telco-grade Kubernetes distribution, and high-performance networking coupled with telecom-specific automation and service assurance.

Telco Cloud Platform helps CSPs expedite the innovation cycle to deliver new applications and services on a horizontal platform, reduce operational complexities, and achieve substantial TCO savings over other network deployment approaches, which in contrast typically use purpose-built hardware and siloed architectures. The flexibility and agility of Telco Cloud Platform further accelerates the cloud modernization journey to 5G.

6WIND Virtual Security Gateway

The 6WIND virtual security gateway (vSecGW) is a key component of the 6WIND Virtual Service Router (VSR) product suite, offering comprehensive and highly scalable Layer 3 IPsec-based VPN connectivity. Designed to integrate seamlessly with any network, the vSecGW protects VPN connectivity for all types of network traffic in fixed, wireless, or converged environments.

Virtualizing the 4G and 5G Security Gateway

6WIND has successfully completed a demonstration of its vSecGW, in a close collaboration with VMware by Broadcom. This white paper describes the demonstration methodologies and its key performance results.

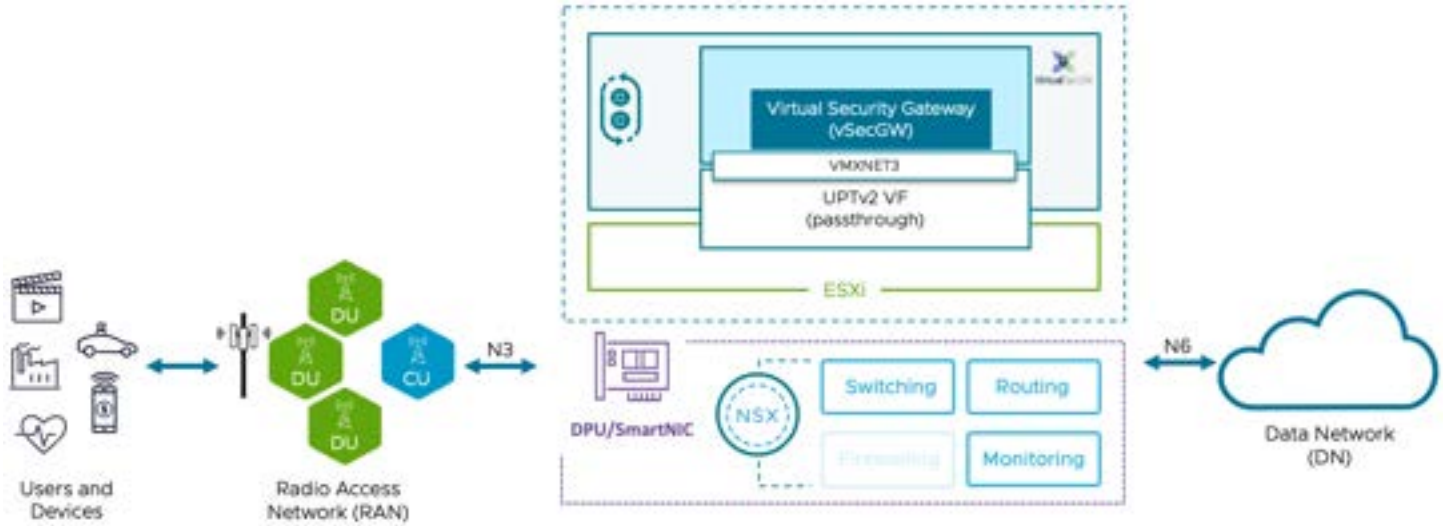

Figure 1: The testing aims to demonstrate that the 6WIND vSecGW can achieve scalable and aggregated (both upstream and downstream) IPsec performance of at least 25Gbps in a network by leveraging data processing unit (DPU) hardware offloading.

The purpose of the demonstration was to show how the 6WIND vSecGW can achieve scalable and aggregated (both upstream and downstream) IPsec performance of at least 25Gbps in a network by leveraging data processing unit (DPU) hardware offloading.

The first phase of the demonstration established the routing performance of the 6WIND vSecGW with 6 CPU cores using different packet profiles.

The second phase of the demonstration illustrated the ciphering performance of the vSecGW using the same CPU assignment and packet profiles as the first phase of the demonstration.

Figure 2: The BlueField-2 DPU - 2x 25Gb/s HHHL form factor.

NVIDIA BlueField-2 DPU

NVIDIA BlueField DPU reinvents the data center architecture by offloading, accelerating and isolating data center infrastructure tasks such as networking, storage, security, and management.

The BlueField-2 DPU is the 2nd generation of SmartNIC from NVIDIA utilizing DPU technologies with a ConnectX-6 Dx networking controller, an array of 16 x A78 ARM processors, DRAM controller, PCIe (PCI Express) switch, and a variety of hardware accelerators in a custom System-on-a-Chip (SoC) installed in each server.

The BlueField-2 DPU offers data plane and control plane offloads to the main server's CPUs. For example, all the network traffic processing or the security tasks can be accelerated and isolated from the virtual machines (VMs) that are running on the CPUs when running on the DPU, delivering line-rate performance, lower latency and higher power efficiency with less CPU overhead.

vSphere Distributed Services Engine on BlueField DPUs

VMware by Broadcom is a leading provider of telco cloud services for applications, enabling digital innovation with carrier-grade control. With vSphere 8, VMware Telco Cloud Platform brings the benefits of cloud to on-premises network functions, supercharges performance through DPUs, and accelerates innovation with a Kubernetes containers-as-a-service solution.

With traditional virtualization technologies, applications require more complex virtual machines with higher networking throughput, resulting in an increase in CPU resources allocated to processing data packets.

Tasks associated with network routing, network overlays (for Geneve or VXLAN), security such as a distributed firewall, telemetry, storage, and remote management can consume more CPU resources on each virtualized server.

One way to reduce this CPU resource overhead is to offload specific tasks, such as networking and security, from the CPU to a DPU, which contains purpose-built engines for handling such tasks.

In 2022, VMware introduced the 8.0 version of its vSphere with new Distributed Services Engine (DSE) capabilities to offload VMware NSX networking and security services onto a DPU, such as the NVIDIA BlueField DPU.

vSphere 8 includes a DSE feature that enables running a custom version of VMware ESXi™ on the DPU ARM processors. ESXi instances running on the DPU and main CPU communicate through trusted interfaces built on top of PCIe. However, all this complexity is hidden from the CSPs and their network functions, to which the DPU works like a regular network interface card (NIC) attached to the server. The installation of ESXi to both the main CPU and DPU ARM environments is seamlessly handled through automation. DPU lifecycle management and monitoring is managed through VMware vCenter®.

This DPU offloading capability frees up CPU cores by running workloads and end-user applications on DPUs while accelerating networking and security performance.

Starting with vSphere 8, VMware introduced Enhanced Data Path Uniform Passthrough version 2 (EDP UPTv2) compatibility, a mode in which a VMXNET Generation 3 (VMXNET3) adapter can be configured to use the capabilities of the DPU in passthrough mode.

Traditionally, a device that is connected in this passthrough mode uses SR-IOV (Single Root IO Virtualization) to bypass the hypervisor to achieve higher networking and security performance; however, this configuration forfeits the use of advanced virtualization features, such as vSphere vMotion and vSphere High Availability, since the flow of data bypasses the hypervisor.

The EDP UPTv2 mode, on the other hand, brings the best of both worlds by enabling network operators to achieve high networking performance without forfeiting the use of advanced workload services that vSphere offers.

Some of this higher performance is because the DPU in EDP UPTv2 mode accelerates network packet processing by performing most of the processing in the DPU hardware, as compared to a standard NIC that performs packet processing tasks in software on the CPU, resulting in slower task completion.

One of the major challenges of using CPUs to handle networking tasks is the cache pollution caused by running hypervisor network processing alongside application logic on the same host CPU cores. When the cache pollution occurs, the frequently used networking and application information often push each other out of the shared cache, reducing the cache hit rate for both. By offloading and isolating network processing to the DPU, the application logic running on the host CPU cores achieves better cache locality while the networking processes use a separate cache on the DPU, often resulting in significant application performance boosts in terms of latency and throughput.

To be able to use the EDP UPTv2 mode, network operators should check the specific VMXNET3 driver versions and ensure full memory reservation for the VMs and availability of DPU Virtual Functions (VF) on the host.

If a network operator does not want to adhere to the restrictions associated with the EDP UPTv2 mode, the DPU can also be used with EDP in the default non-UPT mode, also called emulated mode. This emulated mode also provides DPU-based offloading to accelerate network processing; however, it incurs CPU resource overhead on the CPU host for tagging packets that need to be offloaded. While this emulated mode should also free up some CPU core resources, it will not free up as much CPU core resources as the EDP UPTv2 mode.

Using vSphere vMotion for the EDP UPTv2 mode is achieved by automatically switching the workload migration process to the emulated mode temporarily during the process, and then back to the EDP UPTv2 mode on the destination host upon the process completion, provided the destination also has a Smart NIC VF available. This temporal switching enables network operators to utilize vMotion capabilities and related operations in the EDP UPTv2 mode in a way that appears seamless to them. Note that using EDP with either UPTv2 or the emulated mode requires the NSX Manager to configure networking. Moreover, using EDP UPTv2 may have additional vSphere licensing requirements. Contact a VMware by Broadcom sales representative for details.

vSphere 8 can be deployed on a flat network, where all nodes on the network segment are visible from one another, or with an overlay network tunneling protocol, such as Geneve and VXLAN. Overlay networking creates and manages virtual Layer 2 networks that can span different Layer 3 subnets, allowing a set of nodes to act as if they have their own dedicated network that only they can access while many such virtual networks can co-exist on one physical network. Therefore, overlay networking offers various key benefits, such as tenant separation and enhanced security, in many cloud deployments. Note, however, that not all vSphere deployments require overlay networking.

The remainder of this paper refers to the NSX Enhanced Datapath virtual switch with DPU in full acceleration and offload mode, referred to as Unified Passthrough Mode version 2 (UPTv2).

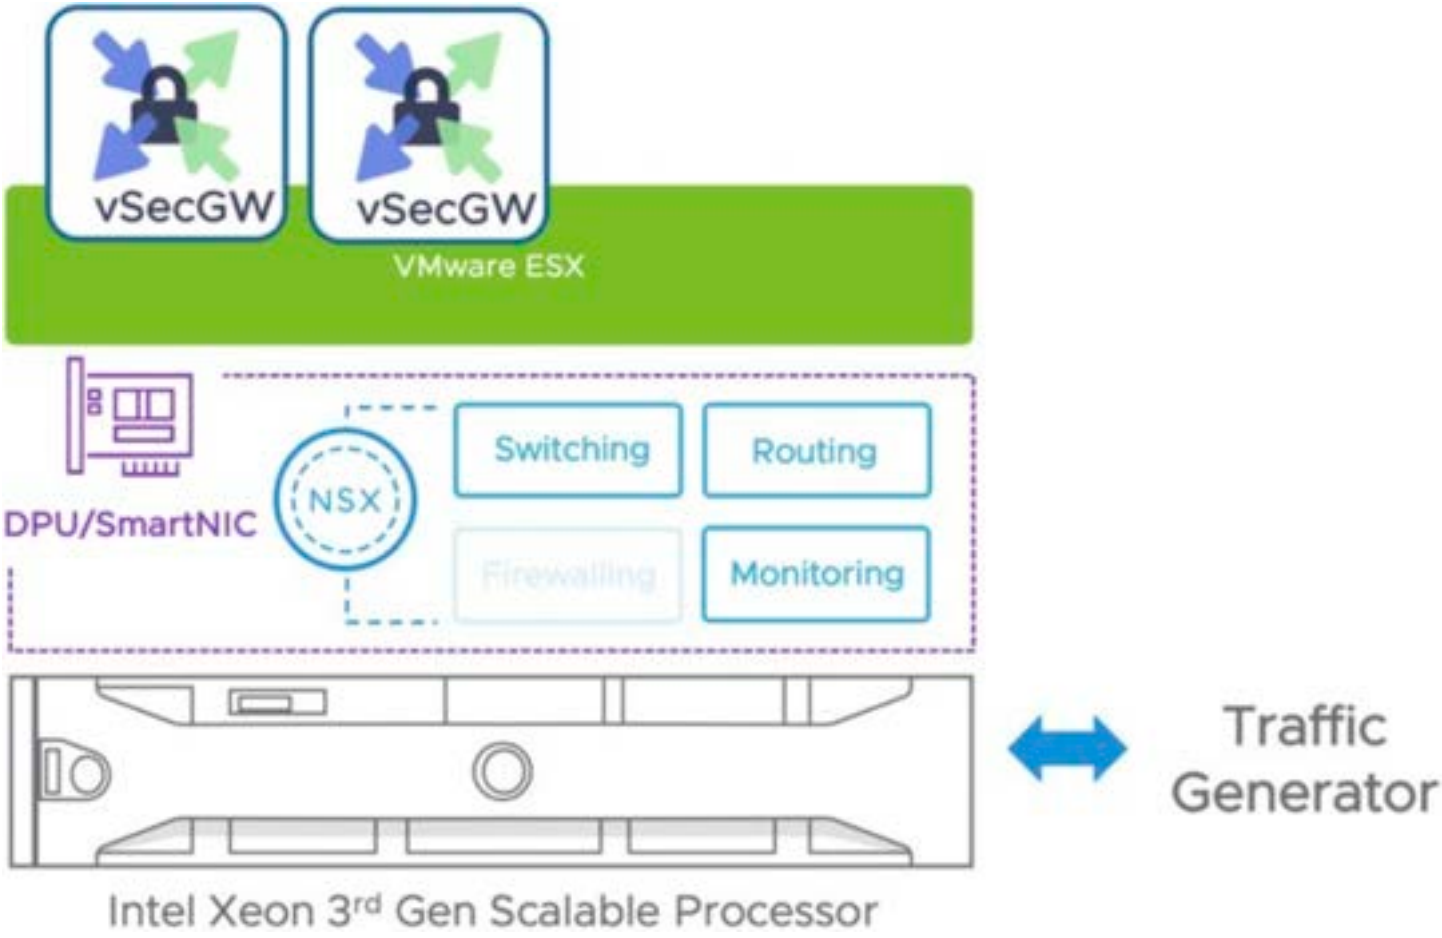

Figure 3: DPU offloading logical topology.

Goals of the Joint Testing

6WIND and VMware have conducted a joint benchmark testing to demonstrate the scalability of 6WIND vSecGW on vSphere 8. Testing was conducted in the VMware labs.

System Configuration

To align with the generally available 4.0 release version of the advanced package of Telco Cloud Platform, the GA version of VMware vCenter Server 8.0 U2c and VMware ESXi 8.0 U2b were installed on commercially available servers while VMware NSX 4.1.2.1 was utilized to manage the virtual switch. A single virtual switch was created by using the NVIDIA BlueField-2 DPU offload capabilities and was being exposed to the ESXi hypervisor.

The testing used PowerEdge R750 servers from Dell with one Intel Xeon Silver 4310 CPU running 24 physical cores at 2.10GHz and 256 GB of main memory.

While Dell servers offered a standard NIC for management, the testing focused solely on the NSX EDP running on the 25GbE NVIDIA BlueField-2 DPU.

Test Setup and Procedures

Figure 3 illustrates the setup for the testing. The test setup includes a traffic generator to emulate typical 4G and 5G User Plane traffic profile.

The test methodology followed the well-accepted Internet Engineering Task Force (IETF) Request for Comments (RFC) 2544 as a guideline.

The benchmarking methodology described in the RFC for throughput measurement specifies a set of packet sizes, which are followed herein. These methodologies relied on test conditions with constant packet sizes, with the goal of understanding the network device capability.

The traffic generator in Figure 4 generated bidirectional traffic that was injected into the DUT. The test traffic packet size distribution, depending on the test, is defined as follows:

• RFC 2544: 256, 350, 512, 650, 700, 1024, 1400, 1500 bytes.

Test Description	
Traffic direction	Bidirectional. Full duplex with traffic transmitting in both directions.
Traffic protocol	UDP/IP
Test run duration	60 seconds
Number of Flows	Single and multiple traffic flows can be generated in the tester by changing the source or destination MAC/IP in the traffic stream.
Acceptable drop rate (also referred to as Packet Drop Rate (PDR)):	0.001%. Binary search algorithm is used to measure the throughput. This algorithm keeps searching for that maximum throughput in packet per second (PPS) contained with the acceptable drop rate.

The test benchmarking methodology used a mixture of packet sizes as well, or Internet mix (IMIX), that simulate telecom operator-specific 5G traffic with the following distribution:

• IMIX1 (average: 350B): 64B (58.3%), 590B (33.3%), 1514B (8.3%)

• IMIX2 (average 700B): 64B (8%), 127B (36%), 255B (11%), 511B (4%), 1024B (2%), 1539 (39%)

The Device Under Test

The DUT consisted of a 6WIND vSecGW running as a virtual network function (VNF) on an Intel x86 Dell Server.

Server description:

• PowerEdge R750

• Intel Xeon Silver 4310 CPU @ 2.10GHz

• NVIDIA BlueField-2 DPU with 2 x 25GbE ports.

• Hardware offload is used.

Hypervisor Description:

• VMware ESXi 8.0.2, 22380479

Virtual Network Function:

• 6WIND vSecGW v3.9 The 6WIND vSecGW VNF can be downloaded through the 6WIND evaluation portal at https://portal.6wind.com/register.php.

The Traffic Generator

The Spirent Test Center (STC) traffic generator was used to generate clear and unencrypted traffic with a pre-determined yet programmable packet-size distribution up to the line rate of the DUT (25 Gbps aggregate).

The traffic generator also emulated the traffic for a User Plane 4G and 5G packet core network, handling a 25G (asymmetric) of raw clear traffic.

The traffic generator in Figure 5 generated bidirectional traffic that was injected into two chained DUT instances simulating an IPSec VPN in software between the two instances.

This testing used the AES-256-GCM algorithm, as it provides the optimal performance while delivering high security.

Test Results

The first phase of the testing was focused on showcasing the performance scalability of 6WIND vSecGW. The different measurements were performed with a vSecGW utilizing different numbers of CPU cores, such as 6 cores.

During this first phase, the traffic generator injected a symmetrical and unencrypted traffic into dual chained vSecGW instances in each direction. Figure 6 shows the test results.

During the second phase of the testing, the traffic generator injected a 25Gbps of symmetrical and unencrypted traffic on a single vSecGW instance in each direction. However, traffic was encrypted with IPSec AES-GCM in software between the two instances.

Figure 7 illustrates that even with IPsec being enabled, depending on the traffic profile (IMIX and IMIX2), line-rate can be achieved with the same number of CPU cores allocated to the vSecGW network function.

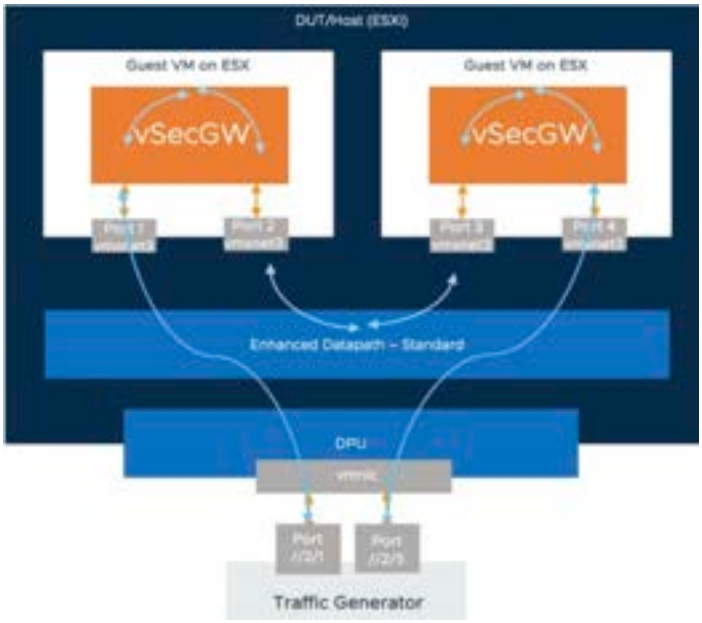

Figure 4: Bidirectional, clear, and unencrypted traffic being injected into two chained DUT instances.

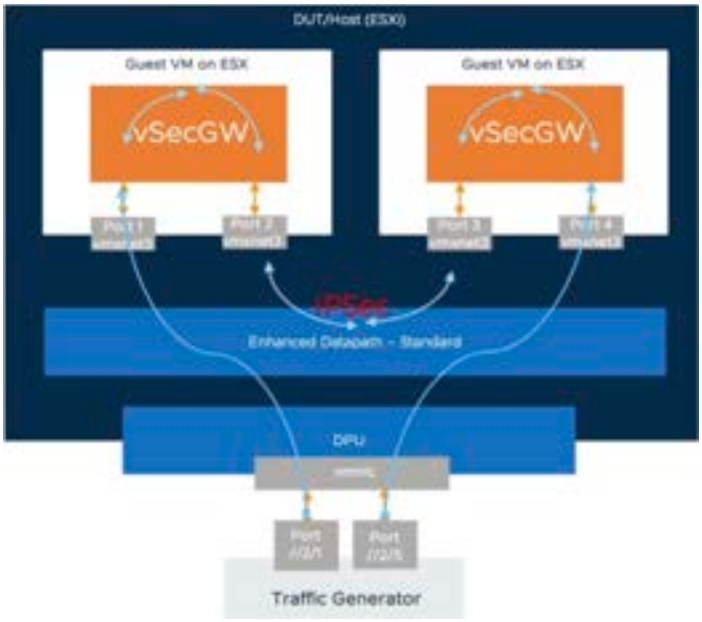

Figure 5: Bidirectional, clear, and unencrypted traffic with IPSec being injected into two chained DUT instances.

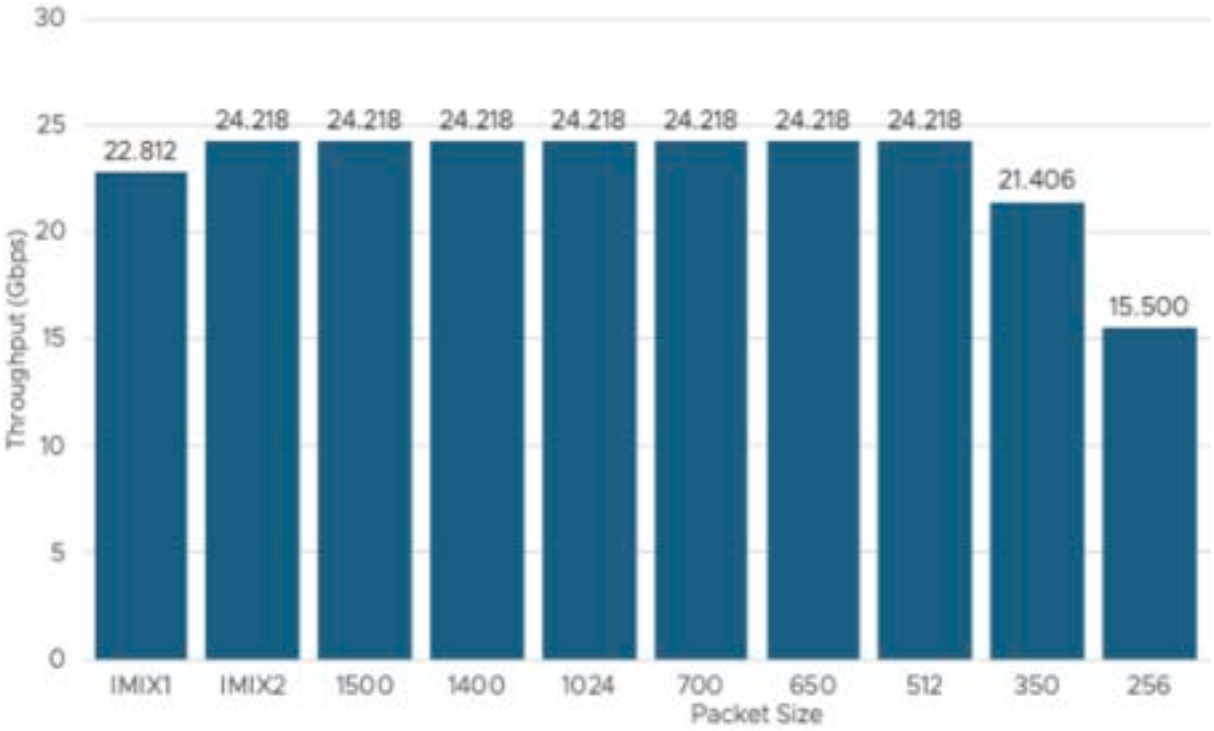

Figure 6: Test results for symmetrical, unencrypted, and NO IPSec-enabled.

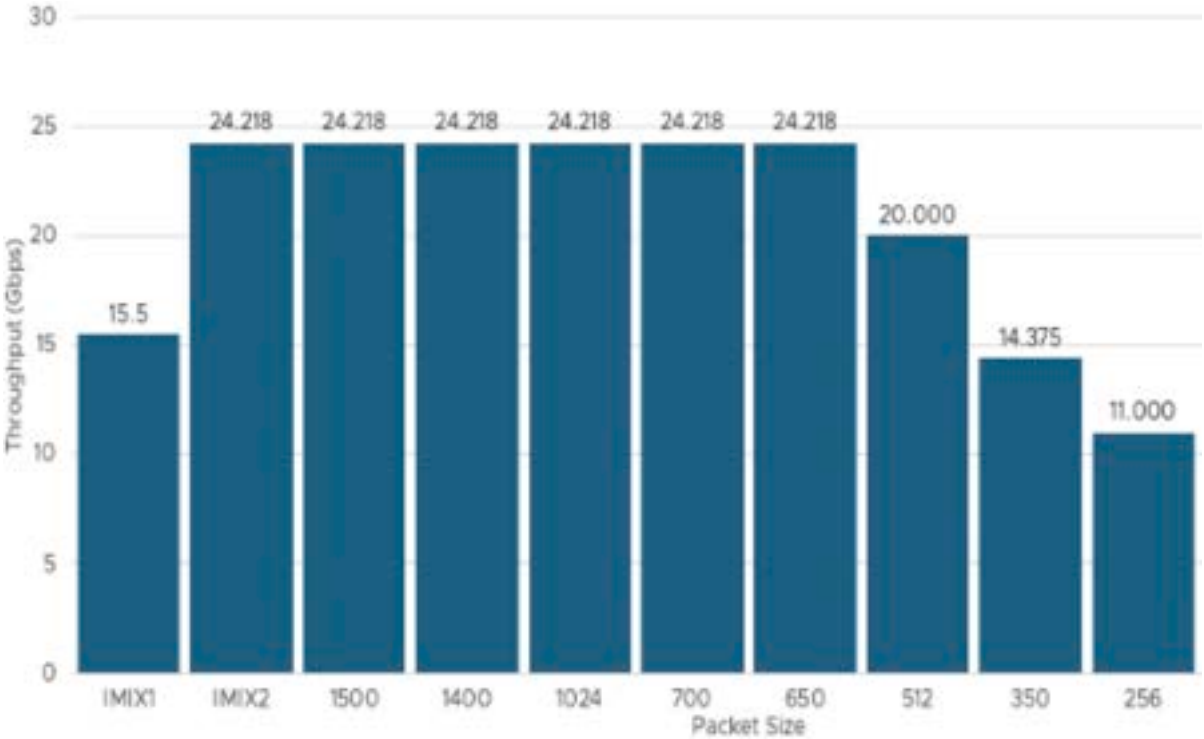

Figure 7: Test results for symmetrical, unencrypted, and IPSec-enabled traffic running through a single vSecGW instance.

Note: The test setup was capped to 25Gbps due to the NVIDIA BlueField-2 DPU with 25GbE being used.

Conclusion

The testing described in this white paper validates that the 6WIND vSecGW, running on VMware Telco Cloud Platform, can efficiently and effectively support IPsec traffic using VMware vSphere with hardware-accelerated networking offloads, supported by NVIDIA BlueField-2 DPU. The test results illustrate that the higher throughput and faster application response times can be achieved regardless of the packet size with an IPsec.

Offloading performance-intensive tasks, such as networking, to the DPU also frees up CPU cores that can be repurposed to support business-critical applications while reducing power consumption and costs due to more efficient practices that can be applied in operating servers and data centers.

In a world facing growing demands for faster packet processing, cloud-native applications, and rising energy costs, all while the demand for energy conservation and green IT are increasing, DPU offloading will become a necessity for both improving application performance and reducing TCO in the data center.

Item No: bc-vmw-telco-performance-dpu-offloading 9/24