

# **EANTC Independent Test Report**

ADVA SmartWAN using 6WIND VSR with Intel®  
QuickAssist Technology Performance Benchmarking

November 2022



## Introduction

The NFV (Network Functions Virtualization) ETSI Industry Specification Group ISG introduced different use cases that represent important service models and fields for the application of NFV. One critical use case was the virtualizing CPE functionalities which show the potential for significant cost reductions by transferring the routing functionality from a purpose-built router to an equivalent capability implemented in COTS hardware.

To combine many fixed-function network components into a single multi-function device, service providers started using Universal Customer Premises Equipment (uCPE). So the network operators can deliver new features like SD-WAN and virtual routing while boosting network security with session border controllers (SBC) and firewall functionality by using virtualized network function software (VNFs) on uCPE.

ADVA and 6WIND created a partnership to optimize uCPE deployments on the Intel® architecture base, including the Dell Virtual Edge Platform (VEP). ADVA's Ensemble Connector provides an operating system for NFV that implements the NFVI software layer of the ETSI NFV architecture. Its purpose is to simplify the deployment and management of uCPEs and support changes in applications or hardware.

6WIND added from their side the routing capability through their VNF, the Virtual Service Router (VSR). According to, ADVA, this software can benefit from a multi-processor environment and optimize data planes' forwarding with its Fast Path feature, which covers essential operations like forwarding, routing, IPsec security, and policing. Additionally, the 6WIND VSR can utilize hardware acceleration when it is available. ADVA has integrated the 6WIND VSR into its Ensemble Connector solution and is marketing this as the SmartWAN vRouter.

## Test Highlights

- Up to 78% throughput performance increase using Intel® QAT for Ensemble SmartWAN vRouter
- Ensemble Connector adds only a little overhead to the scenario. For IMIX traffic, 8% maximum decrease compared to bare-metal in throughput

ADVA commissioned EANTC to perform tests and evaluate the 6WIND VSR performance on uCPE while deploying the Intel® QuickAssist Technology (Intel® QAT) as an acceleration option. Intel developed this technology to offer a hardware solution for cryptographic operations that frees up the CPU for other tasks.

ADVA wanted EANTC to verify the boost the Intel® QAT adds to the performance of the uCPE. Additionally, to demonstrate SmartWAN vRouter's scalability and high network functionality capabilities.

To carry out these tests, EANTC decided to use two different setups. One had a 6WIND VSR loaded on a bare-metal server, while the other had the VNF hosted over ADVA's Ensemble Connector. We wanted the bare-metal platform as a reference point that illustrates ADVA's hosting layer effect.

Then in each setup, we had two rounds of testing: one used software crypto acceleration of the Intel® Multi-Buffer Crypto for IPsec, and the second used Intel® QAT.

We based our test on the main business use case of a universal CPE, an enterprise network with vRouting and VPN functionality, as shown in the following topology.

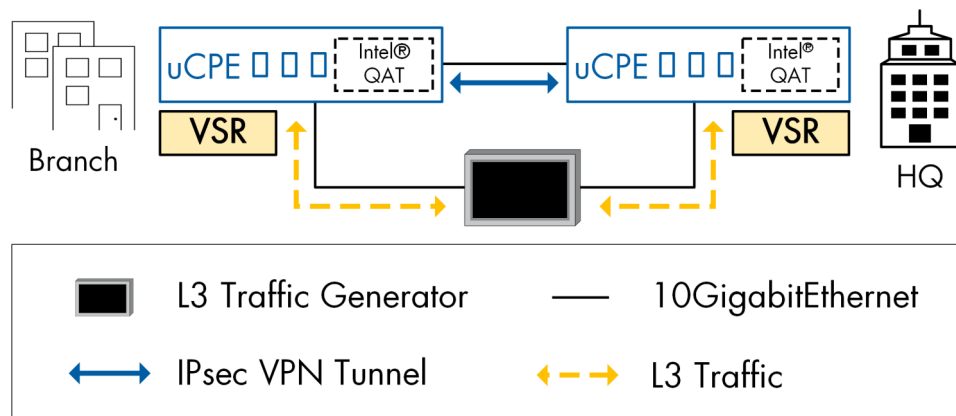


Figure 1: uCPE Business Use Case

EANTC chose to run the test using traffic that simulated real-world traffic patterns and packet distribution. The traffic was dual stack and used internet MIX based on EANTC's actual experiences from enterprise networks, which has an average frame size of 970 bytes with a focus of 38.33 % small packets and 39.99 % large packets. The following table shows the distribution.

Frame Size (Bytes)	Distribution Percentage for IPv4	Distribution Percentage for IPv6
64	5%	-
78	-	5%
100	33.33%	33.33%
373	10%	10%
570	11.66%	11.66%
1256	10%	10%
1518	26.66%	26.66%
9000	3.33%	3.33%

Table 1: IMIX Distribution

Additionally, we generated traffic using single frame sizes. One was 1400 bytes of traffic and a small Frame size (64 bytes for IPv4 and 78 bytes for IPv6), so the ADVA's Ensemble Connector could demonstrate what performance is reachable in high traffic load situations. With large frame sizes, we tested the maximum throughput and, with the small ones, the computing ability.

We installed two uCPEs connected back-to-back for each setup. Then established, an IPsec tunnel between them, over which EANTC delivered unencrypted IPv4 and IPv6 traffic encrypted by the uCPE. When Intel® QAT was enabled, EANTC observed a 95% boost in performance for large packet frames (1400 Bytes) for the bare-metal system and a 75% increase for IMIX traffic. In Ensemble Connector scenario, we had a 90 % and 78 % increase in the throughput for the large packets and IMIX, respectively.

Dell has designed Dell VEP1445 to be optimal for hosting VNFs. It's equipped with an Intel Atom® C3758 processor at 2.2 GHz, which they believe is suitable for low-power requirements and offers a variety of core configurations.

For the 6WIND VSR running over bare-metal, the Intel® 10G ixgbe 5.1.0 driver was used, and an ixgbev 4.1.0 driver for Ensemble Connector hosting scenario.

For testing both bare-metal and hosted VNF, EANTC allocated 3 CPU cores (1 for the control plane + 2 for the data plane) and 8 GB RAM to the 6WIND VSR.



Figure 2: Dell VEP1445

### Testbed Description

The uCPE pair running Ensemble Connector had SR-IOV capable network ports so that we could extend the concept of virtualized functions down to the physical NIC. So Ensemble Connector implements physical function (PF) drivers associated with physical NIC, while the VNF must implement the appropriate virtual function (VF) driver to communicate directly to the NIC.

In our test, we used SR-IOV, which enabled the creation of Virtual Functions from a single Intel® QAT acceleration device to support acceleration for multiple virtual machines. We could generally have 16 Intel® QAT VFs, each attached to a guest Virtual Machine (VM). In our test, we used one Intel® QAT VF.

The Ensemble Connector also enables cloud-native deployments by running OpenStack in Docker containers. This allows benefits from the various functions expected from virtualization/cloud management software.

The physical hardware details are shown in the following table. Figure 3 shows the components.

	uCPE
Server	Dell VEP1445 (210-ASHR)
BIOS	3.48.0.9-4
CPU	Intel Atom® C3758 2.2GHz
Number of Cores	8
NICs for management (onboard)	Intel® i350
NICs for data plane workloads	Intel® X553 10 GbE SFP+ (Firmware: 0x8000084b) (LAN interfaces on the processor)
RAM	16GB DDR4
Disk	240GB SSD
Host OS	6WIND VSR (Bare metal) or Ensemble Connector
QAT Adapter	Co-processor Intel Atom® C3000 Series
VSR Software	6WIND VSR (Virtual Service Router) version 3.4.5-m2
NIC driver (VSR)	ixgbe 5.1.0 (over bare metal), ixgbev 4.1.0 (over Ensemble Connector)
DPDK	21.11
Ensemble Connector Software	Version: 21.1.3
Ensemble Connector NIC driver	ixgbe v5.11.3

Table 3: Physical Hardware

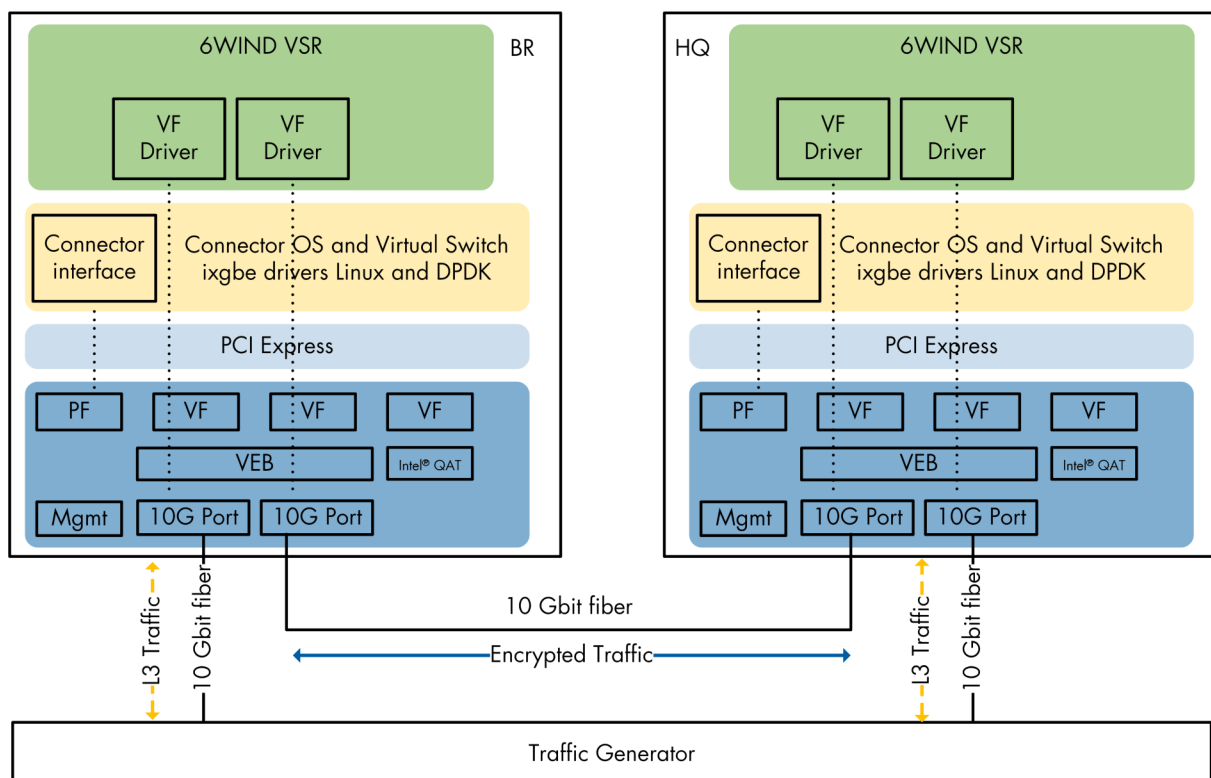


Figure 3: Logical Test Setup

EANTC created the test setup requiring implementing IPsec tunnel between the uCPEs to illustrate Intel® QuickAssist Technology capabilities on the VEP Scalable platform for an NFV-related scenario. The encryption and compression methods employed here were as follows:

### Intel® Multi-Buffer Crypto for IPsec Library

The library primarily targets packet processing applications using software crypto acceleration. It consists of a group of functions that carry out the encryption and authentication algorithms of high computational cost. It takes advantage of the benefits of multi-buffer processing to implement Intel® AES New Instructions (Intel® AES-NI) for encryption.

### Intel® QuickAssist Technology (Intel® QAT)

It's a hardware accelerator that provides symmetric cryptographic acceleration services. Intel® QAT technology is widely available in many x86 servers, and it's used to offload the encryption/decryption and authentication processing tasks to hardware capable of optimizing those functions. This acceleration service releases CPU cycles or up to complete CPU cores. The processing of CPU-intensive workloads can be sped up. Additionally, the available cores can now be utilized for applications of higher importance, such as the operation of virtual network functions (VNFs) and services.

### Test Overview

EANTC used a Layer3 traffic generator to emulate unencrypted traffic with three different traffic patterns to the first device under test, which emulated a branch. The network interface card on that uCPE was the Intel® Ethernet X553 Controller 10 GbE SFP+ 10 Gbps Ethernet ports. The branch used one IPsec tunnel and encrypted the traffic either by hardware acceleration with Intel® QuickAssist Technology or by using the Intel® AES-NI multi-buffer library. The encrypted traffic was forwarded by the uCPE via an Intel® X553 to the second device under test, which emulated the headquarter.

This device under test terminated the IPsec tunnel and decrypted the data (again, either using Intel® QAT, or Intel® AES-NI). The resulting unencrypted traffic was returned to the traffic generator for throughput analysis. The test ran bi-directionally to enable each device to over 20 Gbit/s line speed.

The AES\_256\_GCM algorithm was used in the IPsec encryption/decryption, and CPU isolation and pinning were enabled on Ensemble Connector.

The following table gives an overview of the executed test scenarios.

Test #	Platform	Intel® QAT	I/O
1	bare-metal	No	Direct Assignment
2	bare-metal	Yes	Direct Assignment
3	Ensemble Connector	No	SR-IOV
4	Ensemble Connector	Yes	SR-IOV

Table 4: Test Scenarios

### Test Traffic Parameters

We determined the maximum throughput based on RFC 2544 methodology for the forwarding performance test for fixed packet sizes 64, 78, and 1400 bytes. Also, we tested using the Internet traffic mix (IMIX), which emulates real traffic.

The traffic was bi-directional and mixed evenly between IPv4 and IPv6 (50:50) through 1,000 streams. EANTC considered this a reasonable emulation of the branch-headquarters use case where multiple addresses would be used on both ends.

The following table lists all of the traffic configurations and tunnels that have been used:

Parameter	Value
Traffic Type	Bi-directional L3, Dual-Stack IPv4+IPv6 (50%:50%)
Traffic Flows	500 IPv4 flows +500 IPv6 flows
Test Duration	120 s
Packet Loss Tolerance	10 <sup>-6</sup> (0.0001%)
Frame Size	IMIX, Large Frames, Small Frames <sup>1</sup>
Encryption Algorithm	AES 256 GCM
Authentication	Pre-shared key
Authentication	Secure Hash Algorithm 256 (SHA-256)

Table 5: Test Traffic Parameters

Testing with these values would give a good indicator of the device limit. With small frames, the device needs to handle more interrupts with each new packet, while the large frames, higher performance values could be achieved.

<sup>1</sup> Large Frames: 1400 bytes, small Frames: 64byte IPv4, 78byte IPv6.

## Physical Hardware Details and Virtualization Parameters

Out of the eight available CPUs, two CPUs were assigned for the data plane, with one additional CPU for the control plane.

6WIND recommended enabling Intel® Virtualization Technology for Directed I/O (Intel® VT-d) for the NIC assignment to a VM. DMA-remapping was also required, which allows hardware-based I/O virtualization technologies such as Single Root I/O Virtualization (SR-IOV). So we enabled Intel® VT-d for all scenarios while enabling SR-IOV for the VSR over Ensemble Connector.

Ethernet NICs must also be assigned to the fast path, the Virtual Service Router component in charge of packet processing, to enable their acceleration.

Both CPU pinning and CPU isolation were enabled in all test scenarios to ensure that resource conflicts were avoided.

vCPUs	3
CPU Pinning	enabled
CPU Isolation	enabled
Hugepage Size	2048 KB
Network Ports	1 – management (virtio/DPDK) 2 – data (SR-IOV) 3 – data (SR-IOV)
PCI Device	Intel® QAT VF (SR-IOV)
RAM	8GB
Disk	20GB
Image	6WIND VSR (Virtual Service Router) version 3.4.5-m2

Table 6: Virtualization Parameters

## Results

In total, for the four test scenarios, sixteen test runs were performed, as shown in the following table. The differences between the test runs were the platform (bare-metal or hosted), Intel® QAT yes/no, and traffic patterns with different frame sizes.

To better understand the performance, we added additional latency tests using only IMIX traffic in all four scenarios. The generated load was 90% of the maximum throughput we achieved in the previous tests.

CPU and memory utilization were observed during all tests, and the allocated CPUs for the data forwarding was expected to have the highest load under a throughput stress test.

The data plane CPUs were used in all tests between 99-100%, as expected.

Because of the sole control plane CPU, the devices were always reachable without any issues, even under high load.

We had significantly low throughput for small packets without enabling the hardware acceleration (1.047 Gbit/s for 6WIND VSR running on bare-metal, 0.94 Gbit/s SmartWAN vRouter / 6WIND VSR running on Ensemble Connector). In total, the number of sent frames in these cases was 1,447,864 and 1,298,334 frames, respectively, during the 120 seconds per test run. That indicated that we hit a hardware limitation with this amount of frames. It is important to note that these limitations would not be met in real network situations because a network won't face only small frames.

Therefore we believe our IMIX results for all scenarios are the closest numbers to what a customer might achieve in their actual network load.

The table of results shows for bare-metal platform maximum throughput of 12.2 Gbit/s using 1400 bytes frame and Intel® QAT acceleration enabled while achieving 11.35 Gbit/s for VSR over Ensemble Connector.

For IMIX traffic, we had 8.452 Gbit/s maximum throughputs (with Intel® QAT) for bare-metal vs. 7.9 Gbit/s throughputs for Ensemble Connector.

We could observe that Intel® QAT has an impact, as shown in the Results Table 7.

Intel® QAT	Traffic Pattern	Maximum Throughput Gbit/s		Maximum Latency us	
		Bare-metal	Ensemble Connector	Bare-metal	Ensemble Connector
No	IMIX	4.82	4.420	-	-
	Small Packets	1.047	0.94	-	-
	Large Packets	6.252	5.954	-	-
	IMIX	4.338	3.978	1,263	1,224
Yes	IMIX	8.452	7.900	-	-
	Small Packets	0.905	0.785	-	-
	Large Packets	12.2	11.350	-	-
	IMIX	7.60	7.108	780.28	756.64

Table 7: Test Results



## Intel® QAT Impact on the Performance of the VSR

The following figures show a higher throughput value without packet loss using the Intel® QAT in the case of IMIX and large frames.

For both scenarios, we observed an increment of the forwarding throughput that ranged between 75% and 95% compared to the tests without Intel® QAT . As for small packets with the overhead (IPsec header, padding, ESP header, and trailer) of the offload operation incurring more times per unit of data, we observed decreases in the maximum throughput of the Intel® QAT compared to the software acceleration.

This overhead could go over a 120% increase per data unit for the small packets, while larger packets reduce the overall overhead and achieve higher physical throughput.

Redirecting very small packets to the accelerator is inefficient since it would require more time than encrypting and decoding using the main CPU.

The table below shows the results from the figures in percentage to show the impact of Intel® QAT .

The increase in throughput with and without hardware acceleration for 6WIND VSR deployed either on bare-metal or with Ensemble Connector

	IMIX	64+78	1400
Bare-metal	75.35 %	-13.64 %	95.13 %
Ensemble Connector	78.73 %	-16.48 %	90.75 %

Table 8: Intel® QAT Impact in Percentage

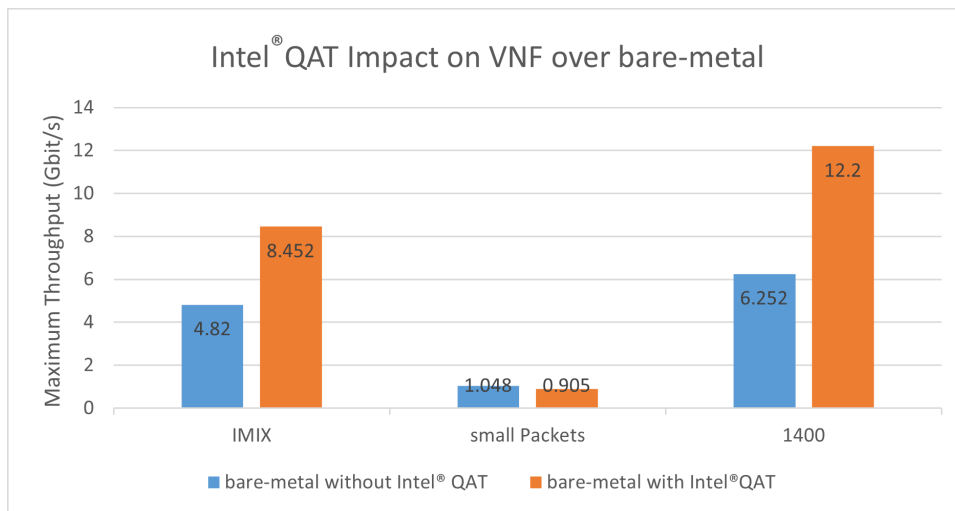


Figure 4-1: Comparison of bare-metal throughput with and without Intel® QAT

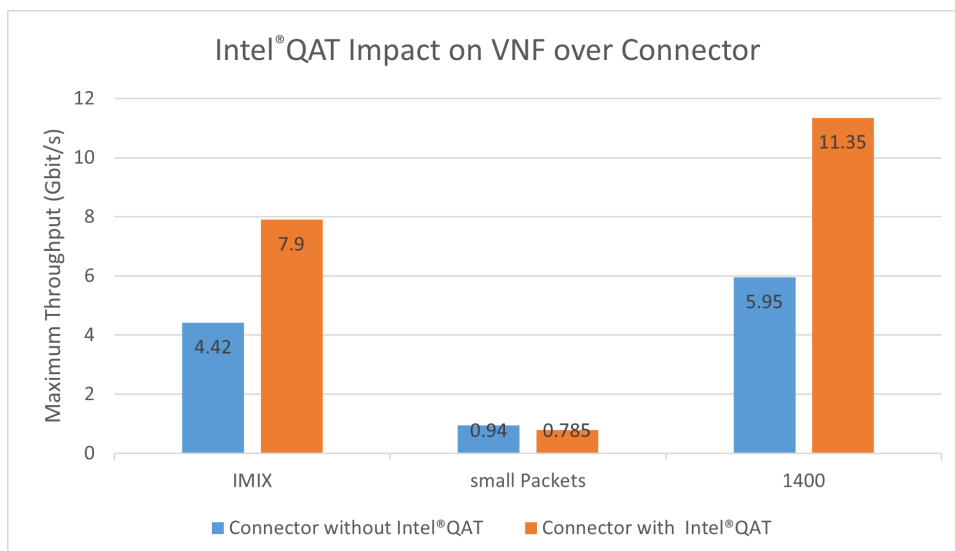


Figure 4-2: Comparison of Ensemble Connector throughput with and without Intel® QAT



## Ensemble Connector Impact on the Performance of VSR

As expected, the performance of the network function hosted directly on the server is slightly better than when running in a VM, with or without Intel® QAT . The reason is that the dedicated resources of interfaces, storage, processing power, and the direct handling of the hardware increased the efficiency and performance of the operating system.

Nevertheless, the difference is only around 9% for the most realistic traffic loads (IMIX and big packets). Enjoying the benefits of virtualization (management, maintenance, service isolation) has only a minimal impact on performance.

We demonstrated this impact in percentage for all scenarios in the following table.

VSR performance hosted on Ensemble Connector relative to VSR operating on bare-metal			
	IMIX	64+78	1400
bare-metal	-8.29 %	-10.30 %	-4.83 %
Ensemble Connector	-6.53 %	-13.25 %	-6.96 %

Table 9: Ensemble Connector Impact in Percentage

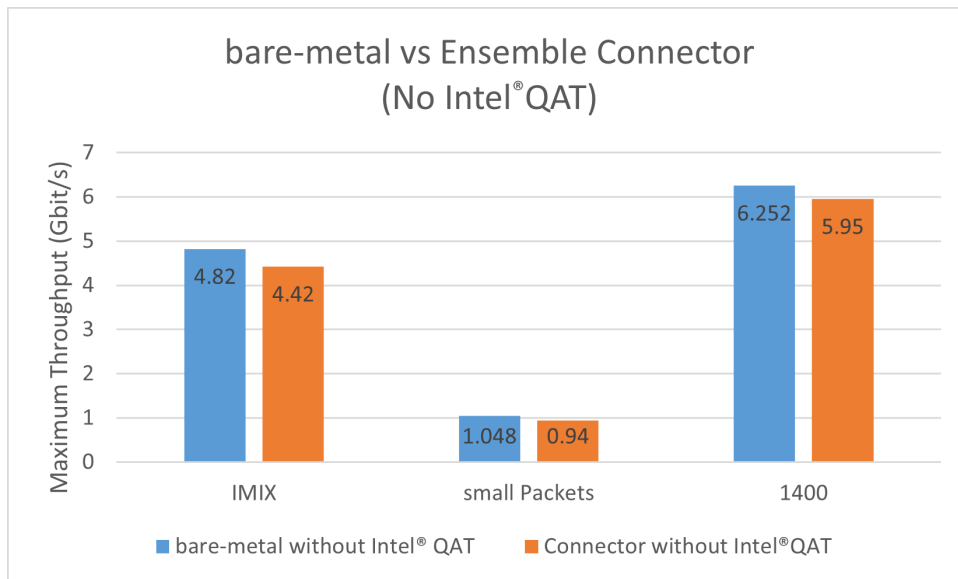


Figure 5-1: Comparison Between Platforms' Throughput without Intel® QAT

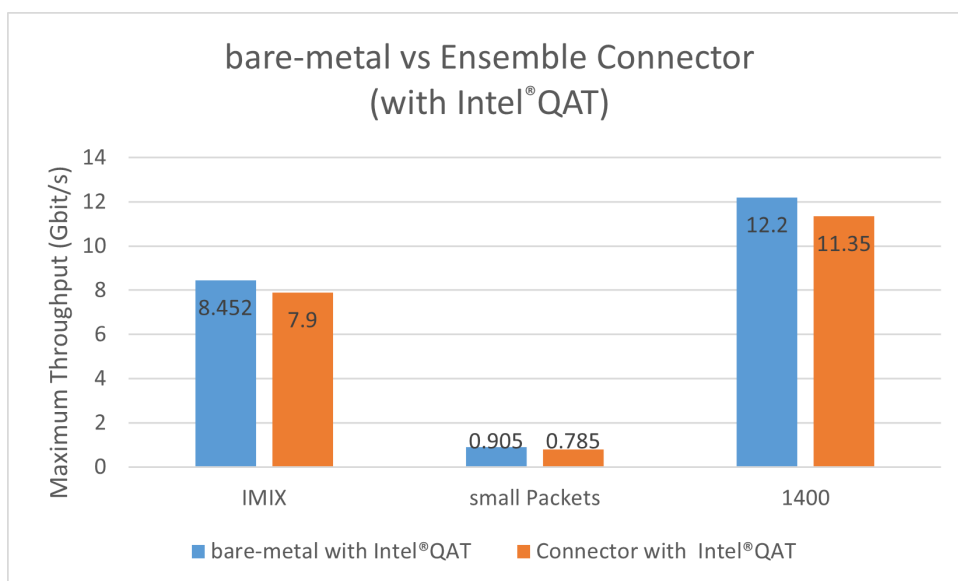


Figure 5-2: Comparison Between Platforms' Throughput with Intel® QAT

## Conclusion

### **Intel® QuickAssist Technology provides a 90% Performance Boost**

EANTC tested the performance of Ensemble Connector hosting the 6WIND VSR. This testing verified that using the hosting server's Intel® QAT to offload encryption processing gave a 78% throughput improvement for IMIX and a 90% throughput improvement for large frames. This improvement was accomplished without increasing the number of CPUs allocated to Ensemble Connector or the 6WIND VSR. Utilizing Intel® QAT will likely reduce the Total Cost of Ownership (TCO) for high-throughput security applications with large packet sizes by requiring a smaller number of CPUs.

### **Ensemble Connector Performance within 10% of Bare-metal**

Additionally, EANTC confirmed that Ensemble Connector reduced throughput by only 10% for IMIX and 5% for large frames. This small impact means that operators can get the benefits of a virtualization layer with minimal impact on performance.

### **6WIND Virtual Service Router (VSR) enables High Performance with Minimal Hardware Resources**

EANTC confirmed the 6WIND VSR is capable of leveraging hardware and software acceleration (when available) to reduce CPU consumption while delivering high performance. On an Intel Atom® C3758 processor using only 2 vCPUs for data plane processing, the 6WIND VSR can deliver up to 6.2 Gbps of IPsec traffic with pure software acceleration and up to 12.2Gbps of IPsec traffic when leveraging hardware acceleration.



This report is copyright © 2022 EANTC AG.

While every reasonable effort has been made to ensure accuracy and completeness of this publication, the authors assume no responsibility for the use of any information contained herein. All brand names and logos mentioned here are registered trademarks of their respective companies.

EANTC AG  
Salzufer 14, 10587 Berlin, Germany  
info@eantc.de, <https://www.eantc.de/>  
[v1.1 20221116]