IPsec Performance Boosts with Networking Platforms based on Intel Xeon Scalable processors

Securing the privacy of mobile subscriber traffic in an increasingly heterogeneous front haul network, and safeguarding the security of the backhaul infrastructure is a never-ending challenge for mobile network operators, as the number of subscribers and devices accessing the network at higher bandwidths continues to rise. Threats will certainly not come to an end when 5G broadband is deployed.

As worldwide security concerns increase, so will the use of VPNs, encrypting, encapsulating and tunneling traffic for authenticated remote and mobile workers, bridging LANs across multiple company sites and linking global data centers together as virtual workloads shift geographically to match both virtual machine and user demand.

This shift places an increasing demand on the encryption capabilities and throughput of servers, as open architecture x86 designs become the preferred platform for the next generation of both bare metal and virtual function VPN gateways, routers or firewalls.

This whitepaper investigates the IPsec performance increases introduced by the new Intel Xeon Scalable family with respect to previous generation by performing benchmarks and throughput measurements using an Advantech SKY-8101 high performance server.

IPsec Performance Boosts with Networking Platforms based on Intel Xeon Scalable processors

Introduction

The need for encryption and the secure transmission of data between sender and receiver is fundamental for network infrastructure integrity and user privacy. In a mobile network for example the evolution to LTE brings a more open and flexible, all-IP network into play, making it increasingly important to boost security at the network edge where it is the most vulnerable. Here, IPsec encryption by Security Gateway (Se-GW) functions is vital to protecting front haul traffic between eNodeBs and the Evolved Packet Core (EPC).

Network security concerns will be further exacerbated by heterogeneous networks and network densification as the infrastructure evolves to 5G and the number of small cell base stations connected to the front haul grows exponentially. The ensuing increase in throughput will certainly require hardware acceleration features in order to handle anticipated mobile broadband speeds and a scalable in-line security solution will be needed, one that has the capacity for millions of high bandwidth IPsec tunnels.

The Advantech SKY-8101 is a high performance server with maximum PCIe expandability in a reduced 1RU footprint and was designed with the high port density and hardware acceleration needed to protect core networks by encrypting, authenticating, and authorizing all data packets that pass through it.

Benchmarking

This whitepaper is destined for network builders evaluating new open Intel architecture platforms that can provide the desired levels of Se-GW functionality and throughput required to protect next generation network infrastructure.

It describes a series of benchmarks which were performed to measure encryption throughput gains when deploying a short-depth, Advantech SKY-8101 Carrier Grade Server based on the Intel Xeon Processor Scalable Family.

The benchmarks described here were selected for their ability to demonstrate performance gains using

(1) the latest Intel AES-NI (Advanced Encryption Standard – New Instructions), an instruction set extension that contains instructions specifically developed for facilitating optimized AES implementations, and

(2) IPsec offload on QuickAssist Acceleration Technology (QAT).

(3) Tests were performed in both virtualized and non-virtualized set-ups. The configurations were then optimized for small and large size packets for further performance analysis.

Device Under Test – Advantech SKY-8101

The key features of the Advantech SKY-8101 device under test (DUT) used for the benchmarks are listed below:

- Single Intel® Xeon® Platinum 8176 Processor with 28 Cores at a base frequency of 2.1 GHz.
- 6 x 16GB DDR4 2400MHz RDIMMs, 6 channels for a total of 96GB
- 2x Intel XVV710 PCI Express Gen3 x8 Network Interface Adapters with two 25GbE Controllers (2 x 25G are used from each card)
- 4x Intel X710 PCI Express Gen3 x8 Network Interface Adapters with quad port 10GbE Controller (4 x 10G are used from each card)
- 8 cache ways allocated for DDIO. SW

Key software features for the benchmarks are listed here:

- Data Plane Developer Kit - DPDK 17.05
- QuickAssist QAT1.7-1.0.1
- Cryptodev Poll Mode Driver (PMD).
- Cryptodev Scheduler PMD (for hybrid IPsec)
- IPsec-secgw
- FD.io VPP IPsec

Diagram 1 below shows the example packet flow for the benchmarks with QAT and Ethernet Adapter in situ. A Spirent hardware test and simulation platform was used to generate packet traffic to the DUT ports and determine the throughput at the tester side. The Spirent was set up as follows:

Throughput Test, Start Traffic Delay: 10 sec, Latency Type: LIFO, Uncheck “Enable Learning”, Trial Duration: 20 sec, Initial rate: 100%

- 2 Raw streams from one port for running RFC2544 with 0.01% acceptable frame loss.
- Tx Port #3, Raw stream, 64B, Dst IP=192.168.105.1, Rx port=Port #2
- Tx Port #3, Raw stream, 64B, Dst IP=192.168.106.1, Rx port=Port #3

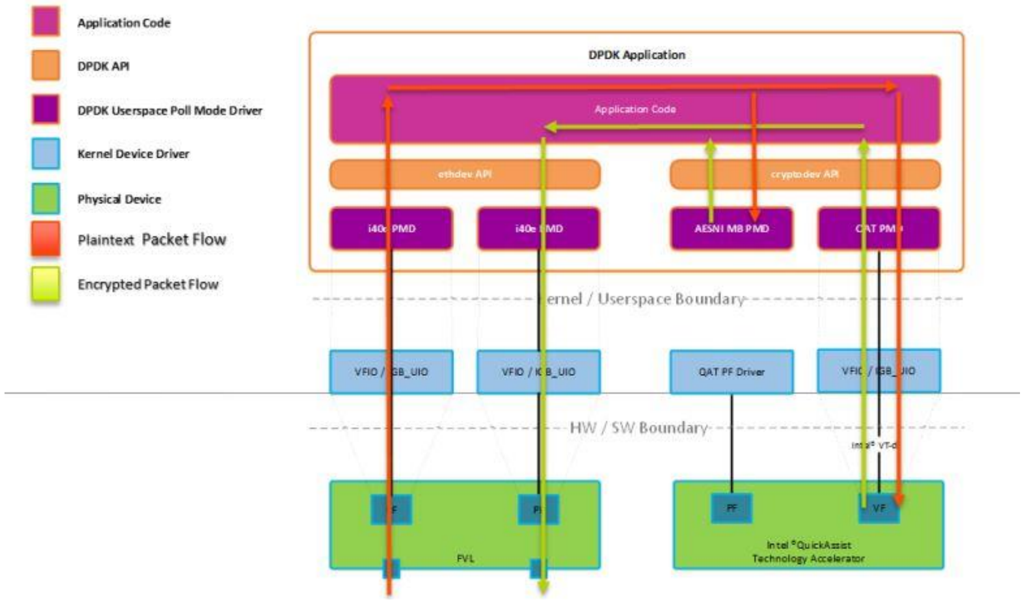

Figure 1 Example Packet Flow for a trivial packet encryption application. Source: https://dpdksummit.com/Archive/pdf/2016Userspace/Day01-Session06-Userspace2016.pdf

AES-128 & AES-256 – Baseline Tests

Initial baseline performance tests were executed to ensure the Advantech platform was set up correctly and matched expected Intel platform performance levels. AES128-CBC, AS128-CBC HMAC-SHA1 and AES256-CBC HMAC-SHA2-256 tests were performed without IO to test the QAT crypto performance.

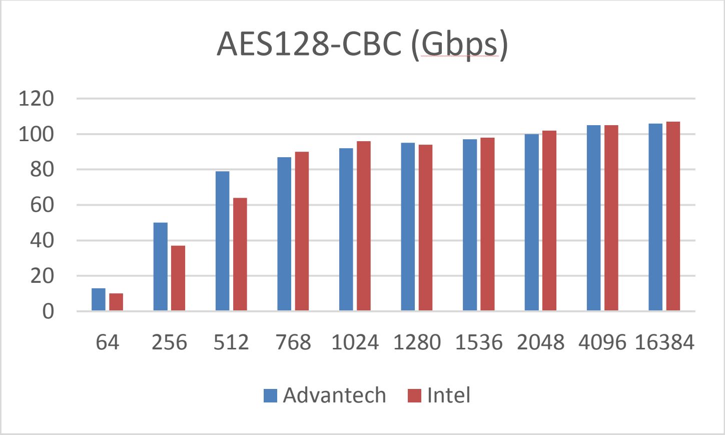

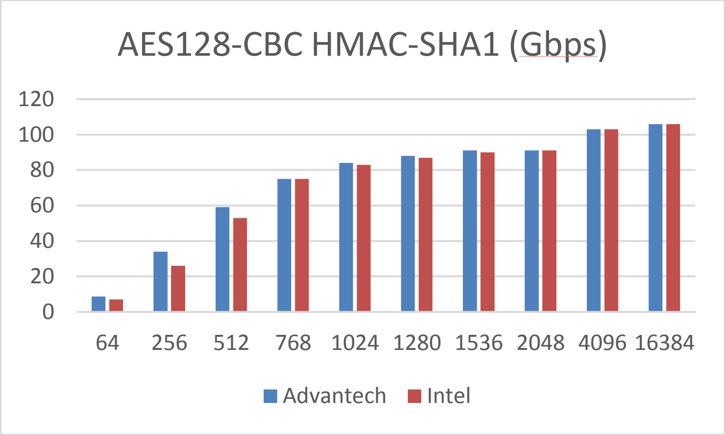

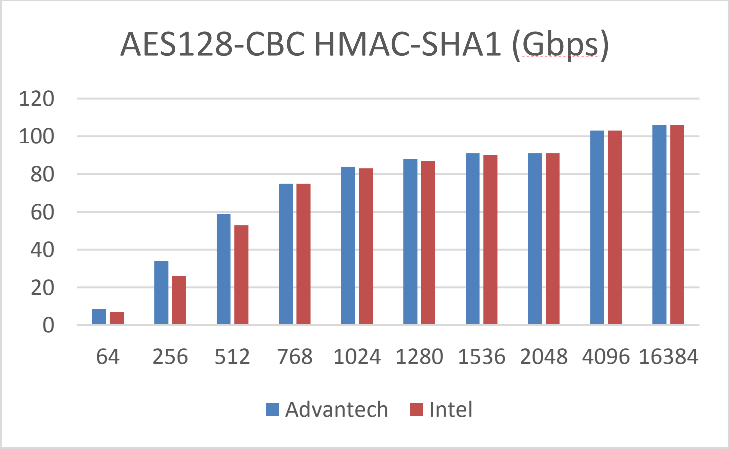

In all cases the Advantech SKY-8101 performed as well as the Intel reference platform and demonstrated better crypto throughput on smaller packet sizes.

AES-NI Tests

The next set of tests compare AVX256 “mb(0.44)-SKL” with AVX512 “mb(0.45)-SKL” on the DUT against a previous generation Intel® Xeon® E5-2658v4 “mb(0.45)-SKL” platform.

A significant performance increase of up to 37% was noted with AVX512 compared with Intel® Xeon® E5-2658v4 throughput results and an increase of up to 29.5% better was observed on smaller packet sizes.

AES128-CBC SHA1_HMAC(Gbps)

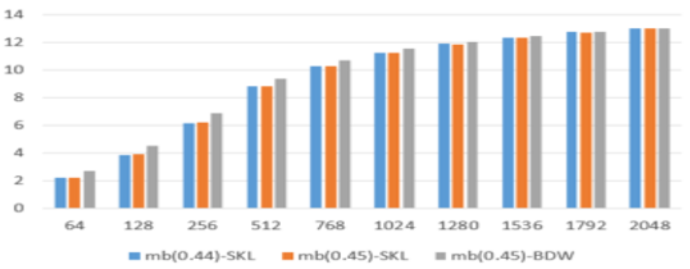

AES128-CBC SHA256_HMAC(Gbps)

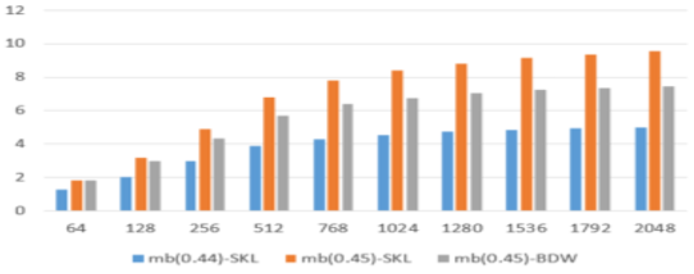

AES128-CBC SHA512_HMAC(Gbps)

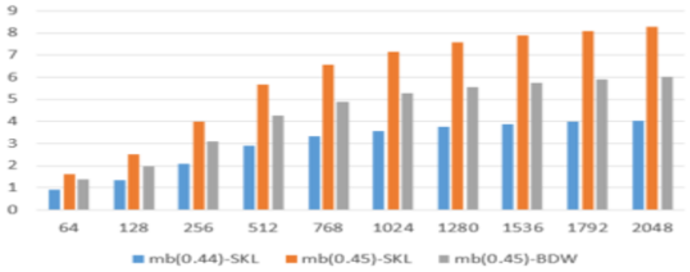

AES256-CBC SHA512_HMAC(Gbps)

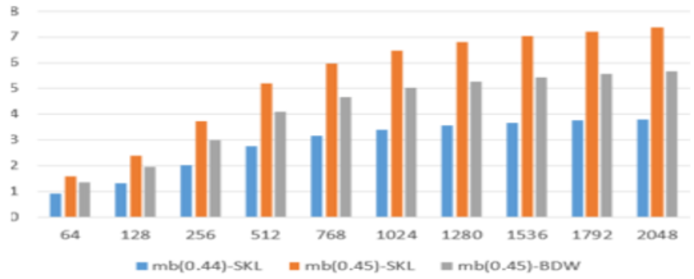

DPDK IPsec-secgw

With the data on AES-NI/QAT as a base line result, the IPsec-secgw test was performed to provide us actual data plane with IPsec throughput based on real network workloads with

- Single core AES-NI
- Dual core AES-NI
- QAT.

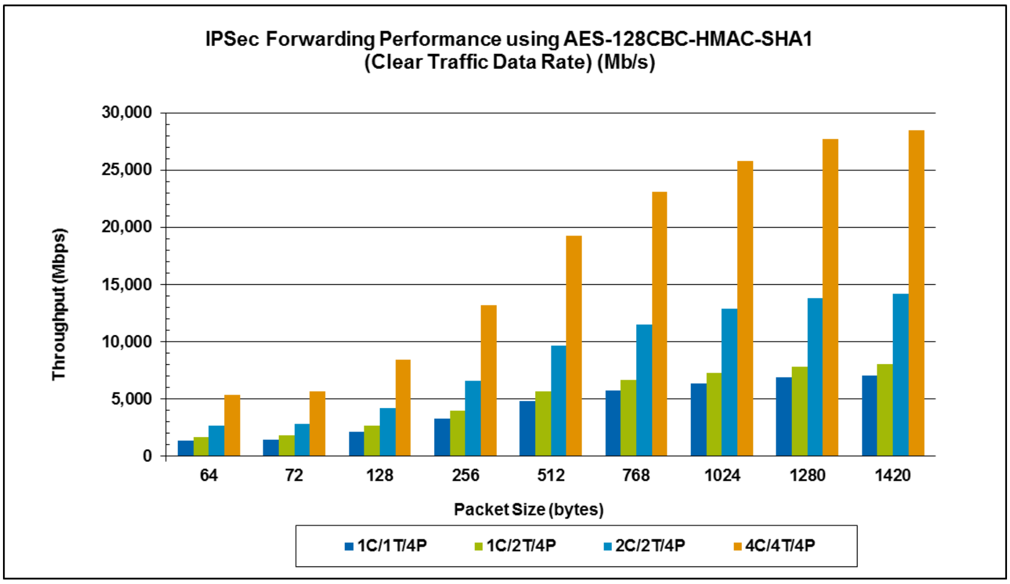

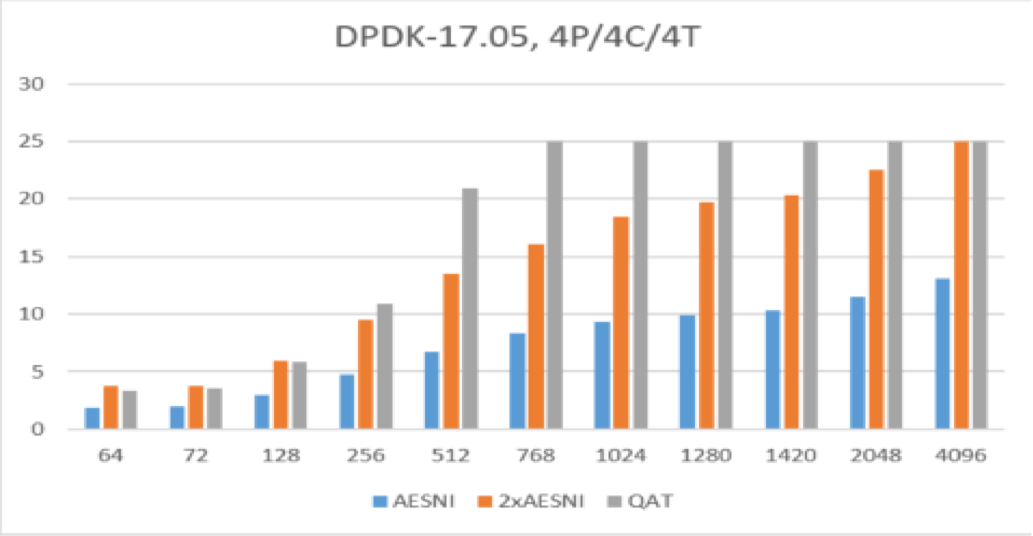

y-axis = Throughput in gigabits per second, x-axis = packet size

The test result concludes that QAT performs better than CPU cores in IPsec forwarding performance.

In addition a comparison is made with an Intel® Xeon® E5-2658v4 platform using DPDK 16.07 AES-NI Forwarding Performance - 4 Ports 1:1 Nodes (bi-directional flows).

Advantech SKY-8101 Series

The SKY-8101, available in both carrier grade and industrial versions, meets_market demands for higher performance, broader scalability, and increased security at the network edge where a new breed of security gateways are required and where new technologies such as virtual Radio Access Networks (vRAN), Edge Cloud, Fog and Multi-access Edge Computing (MEC) are vital to enabling the next generation of digital services.

Both the SKY-8101, and the SKY-8101L with high-capacity storage, also meet the needs of industrial applications where cost efficient, compact, rugged and reliable solutions are required in environments with limited space, higher ambient temperature and low noise constraints.

The Advantech SKY-8101 and SKY-8101L are highly configurable high performance servers designed to balance the best in x86 server-class processing with maximum I/O and offload density in a 1U compact chassis.

The systems are cost effective, robust platforms optimized for superior reliability in business critical applications such as communications, edge and industrial computing.

They are specifically designed for high density PCIe card payloads where maximum I/O connectivity is needed or the integration of industry leading offload and acceleration technology is essential.

The power and cooling options along with the streamlined mechanical design make them ideal for demanding applications requiring high performance acceleration technologies such as GPU, DSP and FPGA cards.

In addition, the SKY-8101L expands storage capacity with support for Intel® VROC hybrid NVMe and SATA RAID that can be leveraged for video caching, data acquisition and storage as well as accelerated edge processing and analytics.

IP Security (IPSec)

IP Security (IPSec) is a suite of security protocols that operates at layer 3 in the TCP/IP layering model. It provides security functionality in the form of confidentiality and authentication for the IPv4 and IPv6 layers. IPSec operates at layer 3, therefore it can provide this protection to all higher level layer traffic (including application traffic) that traverses the internet.

In Linux*, the native 2.6 kernel IPSec stack is called Netkey. It integrates with the Transformer module (XFRM) in the kernel. Netkey accesses the Security Policy Database (SPDB) and the Security Association Database (SADB) to retrieve IPSec policies and IPSec security associations. A user space application, typically an Internet Key Exchange (IKE) stack, is responsible for loading the kernel SPDB and SADB with information necessary for the kernel to establish an IPSec connection.

IPSec Modes

IPSec has two modes of operation, tunnel mode and transport mode.

Tunnel mode is typically used to create a Virtual Private Network (VPN). An IPSec VPN can support secure network-to-network communications, host-to-network and also host-to-host configurations. Network-to-network VPNs are typically used to secure communication between sites. Host-to-network VPNs are often used by remote users that need to connect securely to a corporate network. Tunnel mode VPNs can also be used to secure host-to-host communication (although transport mode is more commonly used in this scenario).

Tunnel mode secures the entire IP packet and encapsulates it in another IP header specific to the IPSec tunnel endpoints.

Transport mode is typically used to secure host-to-host communication. With transport mode, only the IP packet payload is secured. The original IP source and destination addresses remain unchanged

SKY-8101 & SKY-8101L Features

- Compact 20" deep (SKY-8101) & 27.5" deep (SKY-8101L) 1U rackmount servers
- Single Intel® Xeon® Platinum, Gold, Silver or Bronze Processor
- 6x DIMM sockets support up to 384 GB DDR4 1600/1866/2133/2400/2666 MHz SDRAM (ECC/RDIMM/LRDIMM)
- Rich add-in card support: up to 2x FH/FL PCIe x8 Gen3 slots, 1x LP PCIe x8 Gen3 slot, 1x PCIe x4 Gen3 slot for Advantech Personalization card
- SKY-8101: up to 4x 2.5" hot-swappable HDD/SSD drives 1x M.2 2280 SATA SSD
- SKY-8101L: up to 8 x 2.5" hot-swappable HDD/SSD drives and optional 2x 2.5" NVMe SSD drives
- IPMI 2.0-compliant management with reliability and security enhancements
- Optimized platform design for Industrial and Carrier Grade Robustness

IP Security (IPSec)

IPSec Protocols

IPSec has two protocols, Encapsulating Security Payload and Authenticated Header.

The Encapsulating Security Payload (ESP) protocol in IPSec enables confidentiality, authenticity, and integrity. Encryption or Authentication only schemes are possible but not recommended. In tunnel mode using ESP schemes, the outer, encapsulating IP header is not afforded any protection, but the inner IP header can be fully secured. ESP is identified as protocol number 50 in the outer IP header. This is the only protocol that provides encryption. It makes use of the DES, 3DES, and AES encryption standards.

The Authenticated Header (AH) protocol in IPSec enables authenticity and integrity. It provides the framework for all of the features except data confidentiality. Both AH and ESP use a Hash based message authentication code for the data integrity check using either MD5 or SHA-1.

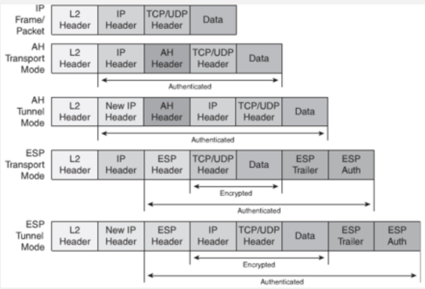

Figure: IPSec Headers

Both AH and ESP work by adding headers to the original packet. Both are transport layer protocols reference by their own IP number

Advantech Contact Information

Hotline Europe: 00-800-248-080 | Hotline USA: 1-800-866-6008

Email: NCG@advantech.com

Regional phone numbers can be found on our website at **http://www.advantech.com/contact/**

www.advantech.com/nc