ETSI NFV Proof of Concept #21
Altera, Broadcom, Huawei, Tilera
Compute Intensive Hardware Acceleration

Introduction

Since the establishment of ETSI's Network Functions Virtualization (NFV) Industry Specification Group (ISG), its stated goal has always been to accelerate the adaptation and acceptance of NFV. The ISG also established a Proof of Concept (PoC) program meant to "build commercial awareness and confidence and encourage development of an open ecosystem by integrating components from different players" (ETSI GS NFV-PER 002 V1.1.1).

Each PoC group formed organically by like-minded vendors and their service provider customers. Our PoC, known by the number the ISG assigns to it (hence PoC#21), included vendors specializing in hardware acceleration boards, CPUs or field-programmable gate arrays (FPGAs) (Altera, Broadcom and Tilera) as well as vendors known for their end-to-end solutions (Huawei).

Background

Together the PoC members wanted to demonstrate the benefits of hardware acceleration in NFV environments for functions that are compute or network intensive. One example of such compute intensive function is IPsec termination. The function involves computationally intensive processes of encryption/decryption and authentication (cryptographic hashing) of the IPsec packet payload. Moreover, processing packets at a high rate and with low latency requires an efficient way to handle communication between the CPU and the network interface.

The PoC members needed a trusted and experienced test lab that could first house all parties and also serve as the central point for coordination, potential mediation between the participants and secondly, a fair and reliable authority on testing and issue resolution should such arise. That role was taken by EANTC. We invited the parties to our Berlin-based lab, created the test plan and executed the tests together with the test vendor – Ixia Communications.

Tested Devices

The participating PoC members provided various solutions to the IPsec gateway function. From a logical perspective all the solutions performed the same function – terminate IPsec tunnels arriving from the emulated clients, decapsulating the traffic and forwarding it to emulated servers. Since bi-directional traffic was used, in the return path, the solutions encrypted the packets generated by the servers and forwarded them back to the clients. The specific products and their function in the test are described below.

Vendor	Product in Test	Function
Altera-+Huawei	Stratix V GX	FPGA-based IPsec acceleration card
Broadcom	XLP832	Universal Crypto-processor
Tilera	TILEncore-Gx	NIC with IPsec acceleration

Test Equipment

The tester we used in the PoC was Ixia's PerfectStorm One. The tester was configured with IKEv1 for authentication type and, when the test profile involved encrypted tunnels, AES-CBC-128 to encrypt the payload. We also verified that the tester performance would not be the limiting factor in the tests by running a loopback test between the ports and verifying that the performance we obtained from the tester exceeded the performance goals planned for the solutions under test.

Methodology

To demonstrate acceleration in terms of IPsec-relevant parameters, one has to be able to measure certain performance using an unaccelerated Network Functions Virtualization Infrastructure (NFVI) (the Commercial Off-the-shelf server) and then use the same NFVI, activate the accelerator function, and repeat the same test. Since each solution used a different approach to acceleration we could compare some functions, such as IPsec session activation rate to the unaccelerated solution provided by Huawei, and in the case of session throughput, simply compare the accelerated performance to the available bandwidth on the system.

In the sections below we provide the results we collected during the IPsec acceleration tests for each participating vendor or ecosystem.

Altera and Huawei

Huawei provided an RH2285v2 server to host the virtual functions for the test. The server is based on two 8-core Intel Xeon E5-2470 CPUs running at 2.3Ghz and equipped with 96GB of memory.

On top of the server, Huawei used Ubuntu Linux (3.2.0.29) as the KVM hypervisor and multiple VMs running Ubuntu Linux (3.2.0.23). Huawei also used a custom version of OpenSwan (2.6.41) with enhancements to utilize the acceleration hardware and provide the IPsec gateway functions.

Initially, we measured sessions activation rate, the function which Huawei and Altera chose to optimize, at 160 session activations per second. At that stage the CPU utilization was measured to be around 38-47%. We made attempts to increase the session activation rate to 240 sessions/sec and detected session failures.

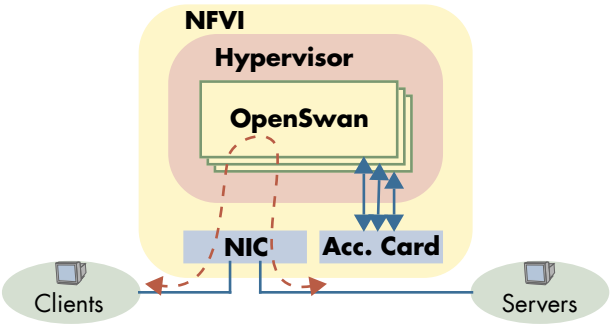

Figure 1: Altera+Huawei Logical Test Setup

In the next step Huawei engineers installed the Altera FPGA card and programmed it using Altera's IPsec accelerator. We repeated the test. The IPsec activation was offloaded to the FPGA card and we could measure a much higher activation rate of 2,000 sessions/sec.

This test demonstrated exactly one of the two approaches to compute intensive functions. By pairing a CPU and FPGA, the heterogeneous hardware remains flexible and ready to implement any virtual network functions. Obviously collaboration between the VNF developer and the FPGA vendor must exist. This type of development is exactly what the ETSI NFV ISG PoCs are meant to stimulate. So it seems that the mission, in this case, is accomplished.

Broadcom

The basis of the solution presented by Broadcom is a single chip with advanced encryption capabilities. The chip provides generic encryption functions and is positioned as a 3rd-party component for other networking equipment solutions. In order to present a testable stand-alone solution for our PoC, Broadcom assembled a XLP832 board equipped with this cryptography chip, a MIPS-based 1.5Ghz central processor, DDR3 memory and two 10-gigabit network interfaces. As Broadcom explained, the XLP series of devices provide a versatile platform for implementing dedicated networking application processing solutions.

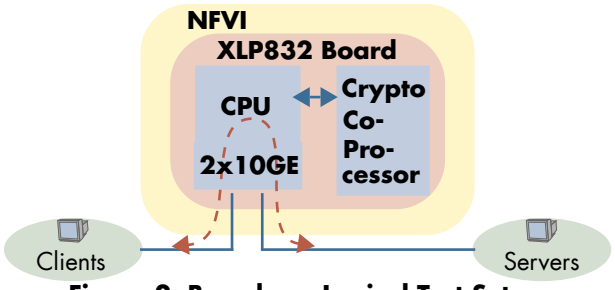

Figure 2: Broadcom Logical Test Setup

The presented solution from Broadcom was designed to accelerate both session setup rate and throughput performance. However, there was nothing for us to compare an unaccelerated performance with – the Broadcom solution was accelerated at all times. We could, however, measure the CPU load of the main processor which was running at 2 – 3% load.

We measured a session setup rate of 4,000 sessions/sec with average session setup rate of 3 ms and maximum session setup rate of 15 – 17 ms. Once the sessions were established, we were able to measure 8.75 Gbit/s per direction throughput performance.

The Broadcom solution demonstrates one additional goal that the PoCs should foster – building commercial awareness. The boards provided by Broadcom are available to both networking equipment manufacturers and to service providers, the latter doing more and more of the integration work themselves to get an edge in the market.

Tilera

Tilera presented a solution that is designed to accelerate the traffic encryption function in an IPsec gateway. The solution comes in a form of a network interface card called TILEncore-Gx equipped with 4 10-Gigabit Ethernet interfaces and a high-performance TILE-Gx 36-core CPU capable of taking over computationally intensive processing of network traffic. The card was installed into a host server using a PCIe interface and in essence was only using power provided by the host.

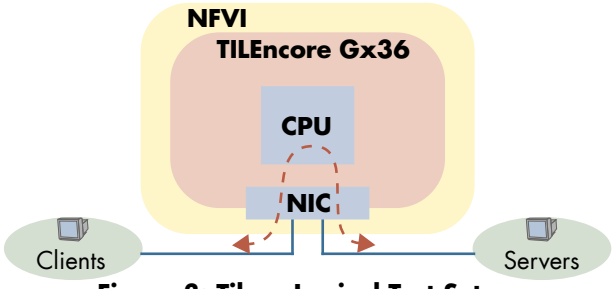

Figure 3: Tilera Logical Test Setup

We used a rate of session establishment of 96 sessions/sec and established 4,000 sessions in total. Over those sessions we generated a total throughput of 29.28 Gbit/s. Currently as Tilera told us, the solution only uses 16 of the 36 cores available and is expected to double the performance in near future. The host CPU was more or less idle.

The impressive throughput rate shown by the Tilera board is the extreme use case of acceleration performance – all IPsec gateway function are done on the board while the COTS x86 is free to perform other virtual functions and in some hopefully near-future case, even run service graphs within the same NFVI.

Summary

In general the PoC shows that we could obtain good results for compute intensive functions, saving on compute resources. The PoC demonstrated nicely two flavors of compute intensive workload offload:

- A function is offloaded – e.g. IKE function is offloaded to an accelerator such as an FPGA NIC.
- A complete VNFC is offloaded to a specialized CPU.

The Huawei setup could be operated with or without the Altera FPGA acceleration card which enabled us to measure two performance values and compare them to each other. The Broadcom solution was an XLP832 board equipped with their newest crypto coprocessor that can serve as a basis for any kind of appliance requiring programmability and high computational performance, while Tilera provided a PCIe card, equipped with its own operating system and capable of operating autonomously.

Given that the solutions presented for the PoC are still in early stages and not yet optimized, as well as, the different approaches that the vendors took, we were not surprised to see different performance values measured for each of the solutions. Nonetheless, all solutions were able to demonstrate the benefits of Hardware Acceleration in NFV environments for compute intensive functions.

About EANTC

The European Advanced Networking Test Center (EANTC) offers vendor-neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, MPLS, Mobile Backhaul, VoIP, Carrier Ethernet, Triple Play, and IP applications.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/

[version v3.0] [20141009], [JG]