Lanner

Network Computing

Innovative Platforms for Secure Networking Infrastructure

Inline Acceleration of VXLAN Tunneling for SDN with Ethernity's SmartNIC

Version: 0.1

Date of Release: 2019-03-19

INTRODUCTION

These days, many companies have outgrown the most basic solutions for meeting the ever-increasing demand for cloud services. For example, virtual local area network (VLAN) usage is limited to 4,096 entities (VLAN IDs), and, given the huge growth in cloud-based networks, much more segmentation is required. A more scalable solution than VLAN is a must.

There are various solutions to address such issues, including multiple VLANs or MAC-in-MAC, which try to overcome the scalability problem by multiplying the number of possible VLAN IDs or MAC addresses that are used. These technologies, though, remain in Layer 2 (L2), which leads to inherent challenges in scaling to an especially large number of entities.

As a result, VMWare, Cisco, and Arista Networks introduced Virtual Extensible LAN (VXLAN) as an approach that successfully addresses this challenge by creating a
virtual network that transfers data across existing Layer 3 (IP) infrastructure. By using tunnels between virtual machines (VMs), a virtual L2 network is overlaid on top of multiple L3 subnetworks. VXLAN's overlay scheme enables cloud-based services to scale without the need to add to or reconfigure existing infrastructure.

However, for the overlay network to work, VXLAN must encapsulate the data so that it can tunnel into Layer 2 and be carried across Layer 3. The VM's L2 traffic is encapsulated with new MAC and IP headers, as well as a VXLAN header that includes a Virtual Network Identifier (VNI), a 24-bit ID that extends the VLAN from 4,096 segments to 16.7 million IDs, which solves the scalability concern.

Yet another advantage to VXLAN arises in the software-defined networking (SDN) architecture. Whereas in a VLAN the need for configuration changes to the existing L2 infrastructure can complicate VM and workload migration from server to server with each server managed separately, VXLAN allows network engineers to migrate VMs, manage and program packet flows in a centralized approach.

VXLAN can be implemented in an SDN scenario over existing hardware and software subsystems without additional investments. However, there are two concerns that arise regarding network performance when using VXLAN:

1. Offloading capabilities of the network interface controller (NIC) are often unavailable because the inner packet is no longer accessible due to the added layer of encapsulation. Thus, additional CPU resources are required to perform tasks that would previously have been handled more efficiently by the existing NIC.
2. Network performance, both in terms of throughput and latency, can easily bottleneck when the number of flows is scaled up, overloading the CPU with packets for decapsulation.

The solution to these concerns is to use a SmartNIC, which can handle offload all VXLAN functions from the CPU to an FPGA, thereby improving throughput, reducing latency, and freeing CPU cycles for other tasks.

LANNER-ETHERNITY VXLAN SERVER-BASED SOLUTION

The Lanner FX-3230 server appliance is designed with a rear expansion of a PCI Express 3.0 X8 (8 GT/s) interface, which is seamlessly compatible with Ethernity Networks' ACE-NIC SmartNIC. The combination of the SmartNIC with the Lanner appliance enables VXLAN offloading to improve performance and save CPU cycles.

To test the solution's ability to handle VXLAN offload, the following test configuration was designed:

- An ACE-NIC was configured for 16K tunnels (although it supports up to 256K in the current generation).
- On the ACE-NIC, 2 x 10G port bi-directional traffic ports were used in this test.
- VXLAN packet size was set to 804 Bytes, simulating an average Internet packet size.
- 4K VXLAN network identifier (VNI) per port direction, equivalent to using 8K streams per port
- Traffic test duration: one minute

Figure 1: Lanner FX-3230 server

APPLIANCE PLATFORM SPECIFICATIONS

Hardware	
CPU	Intel® Xeon™ Platinum 8180 CPU @ 2.50GHz
Memory	32GB DDR4
HDD	512GB MEMXPRO® U.2 PCIe A4E
Software	
Host OS	CentOS 7 64-bit (kernel 3.10.0-693.el7.x86_64)
Hypervisor	QEMU 1.5.3
Guest VM OS	CentOS 7 64-bit (kernel v3.10.0-862.2.3.el7.x86_64)
vSwitch	Open vSwitch (OVS) v2.9.90 with DPDK 17.11.0

Hardware	
(Guest VMs)	

Figure 2: Ethernity Networks ACE-NIC40 SmartNIC

ACE-NIC40 Specifications

SmartNIC	ENET ACE-NIC40 Adaptor Hardware Version: 1302.1.05.1 (Date 12/06/18: hour 20:00)
P/N	ENA1044F
Software version	01.27c MEA Software Version 450.08.102C (May 31 2017 09:25:26)
Host Interface	PCI Express Base Specification 3.0 (8GT/s) X8 lanes
Ethernet Port	4 x SFP + 10G interfaces
FPGA	Xilinx XC7K410

TRex Packet Generator (MB-8896) Specifications

TRex traffic generator software was used in the test to produce stateless uniform UDP packets for test flows. For more information, see the TRex Introduction: https://trex-tgn.cisco.com/trex/doc/trex_manual.html

Hardware	
CPU	Intel® Xeon™ CPU E5-2640 v4 @ 2.40GHz
Memory	32GB DDR4
HDD	240GB Micron® SSD
NIC	Intel XL710 10G NIC with v6.01 firmware (Lanner module: IXM407)
Software	

Hardware	
OS	CentOS 7 64-bit (kernel v3.10.0-862.2.3.el7.x86_64)
TRex	v2.41 with custom Python script for generating VXLAN packets

TEST ENVIRONMENTS AND WORKFLOW

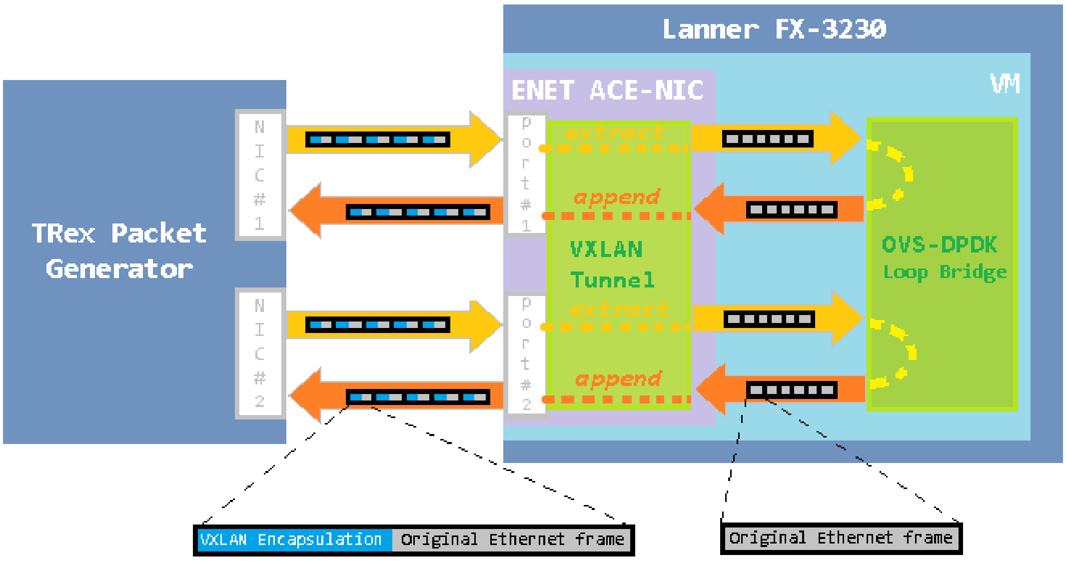

Ethernity's ACE-NIC SmartNIC was selected to be installed onto the Lanner network appliance to test the performance improvement. Ethernity's ACE-NIC is a single standard adapter card with 10G Ethernet connectivity and a programmable FPGA acceleration engine to provide optimal performance, scalability, and cost-effectiveness for NFV, SDN, and VNFs in enterprise data centers.

With support from Ethernity in setting up the VXLAN demonstration, Lanner demonstrated a boost in performance in the test, showing in this scenario 8K streams of traffic with the CPU at 20G line rate. Due to a limitation in the test environment with TRex upon generating more than 8K streams, further experiments were not conducted.

During the experiment, the Ethernity module received VXLAN encapsulated traffic sent from the TRex packet generator software and extracted the VXLAN tunnel so that the internal packet was left to flow towards the CPU.

In addition, the Ethernity ACE-NIC SmartNIC used an OVS bridge with DPDK to provide a loopback flow for offloading tagged VLAN traffic from the CPU. Once the traffic was looped back, the SmartNIC classified this traffic by its DST_IPv4 protocol. It then appended an individual VXLAN tunnel before egressing the traffic from the same physical external ports from which it was received.

CONCLUSION AND TEST RESULTS

The test results for offloading CPU workload and handling VXLAN Tunneling display inline traffic with a high rate speed and zero packet loss in the SmartNIC. Highly available IP packets were forwarded 99.99999% of the time, and any packet loss that occurred in the overall system resulted from DPDK-based transmissions running on a 2-core CPU.

The following screenshot from TRex shows L2 TX and L2 RX at 19.5Gbit/s, indicating line rate transfers and no packet loss.

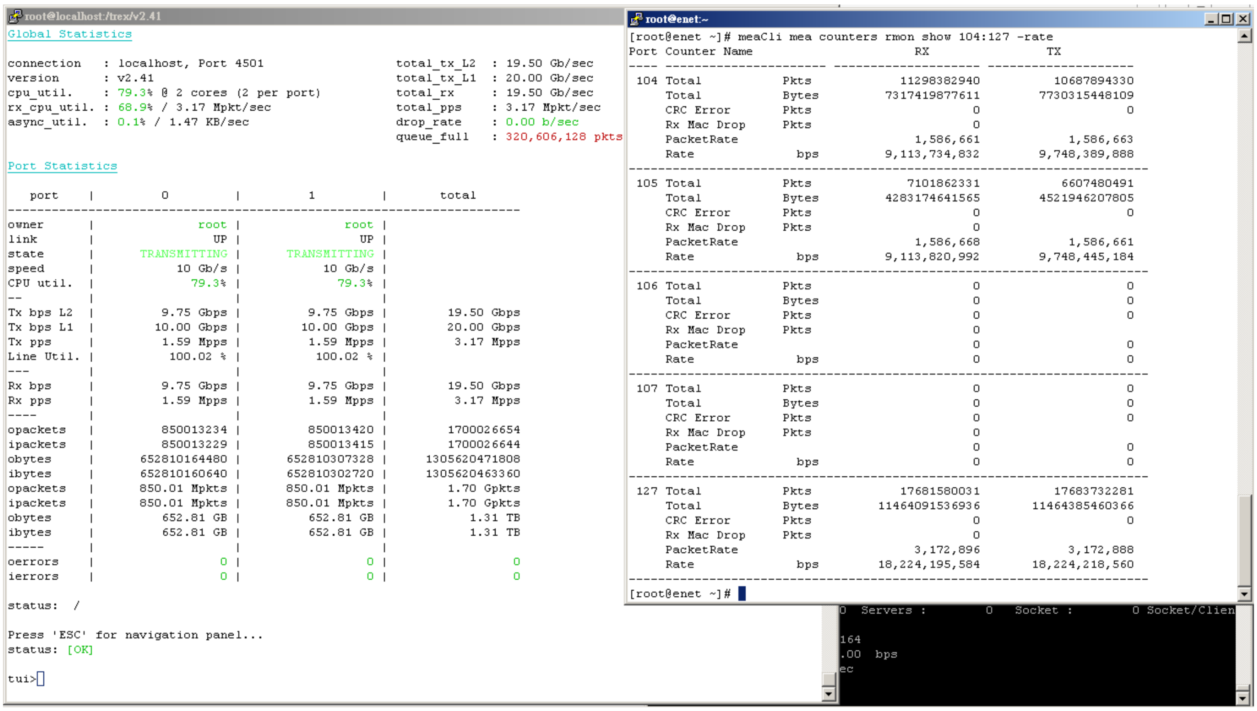

As discussed above, NFV has become a key component in the development of network communications, and the network switch, server, and router have become more critical in NFV network environments for handling technologies like Virtual Tunnel End Point (VTEP) within VXLAN. Ethernity Network's ACE-NIC SmartNICs can optimize and offload CPU workloads more efficiently than handling the processing by CPU alone, especially in network traffic in the Layer 3 and Layer 4 OSI architecture, such as IPSec VPN – essentially enabling the vision of truly programmable networks today.

The Lanner FX-3230 fully complies with Ethernity's ACE-NIC SmartNIC, which provides low latency, high performance, and TOC optimization. The integration of the Lanner network appliance with ACE-NICs offers a powerful customizable, and fast-delivered solution with optimized hardware and state-of-the-art open networking capability.

For more information on this solution, please contact us at contact@lannerinc.com or info@ethernitynet.com.

About Lanner:

Lanner Electronics Inc. (TAIEX 6245) is a world-leading hardware provider in design, engineering, and manufacturing services for advanced network appliances and rugged industrial computers. Lanner provides reliable and customizable computing platforms with high quality and performance. Today, Lanner has a large and dynamic manpower of over 900 well-experienced employees worldwide with the headquarter in Taipei, Taiwan and subsidiaries in the US, Canada, and China.

For more information on the Lanner FX-3230, please see:
http://www.lannerinc.com/products/network-appliances/x86-rackmount-network-appliances/fx-3230

About Ethernity:

Ethernity Networks (AIM: ENET.L) is a leading provider of data processing solutions and technology for high-end Carrier Ethernet applications across the fixed and mobile telecom, security and data center markets. The Company's core technology, which is populated on programmable logic, enables data offloading at the pace of software development, improves performance and reduces power consumption and latency, therefore facilitating the deployment of virtualization of networking functionality.

For more information on the Ethernity ACE-NIC SmartNIC, please see:
http://www.ethernitynet.com/products/smartnics/ace-nic-2/