


vEPC Acceleration Using Agilio[®] SmartNICs



The virtual Evolved Packet Core (vEPC) is the crux of 3G and 4G mobile networks used in service provisioning for mobile users. The main functionality of the vEPC is to act as an interface between the 3G and 4G radio interfaces and public IP networks. The vEPC is implemented as virtual network functions (VNFs) within virtual machines (VMs) running on commercial off-the-shelf (COTS) servers. Netronome Agilio SmartNICs provide a 5X increase in vEPC bandwidth on the same number of CPU cores by using the open vSwitch (OVS) offload and SR-IOV technologies. (See Figure 1 below).

With the emergence of cloud-based server architectures with NFV and SDN technologies, the EPC components can be implemented as the VNFs on the x86 servers while still maintaining high availability and scaling requirements. The implementation of EPC components running on VMs as VNFs has advantages such as hardware independence, scalability and on-demand provisioning.

The vEPC performance measurement with Agilio SmartNICs and traditional NICs is done on a setup as shown in Figure 2. Each of the four servers in this setup is a Dell R730 with 24 physical cores (48 hyper threaded cores or vCPUs on two NUMA nodes). The setup simulates a vEPC network, specifically for GPRS tunneling protocol for user plane (GTP-u) data traffic, the same that is used in the 4G LTE mobile network. In the vEPC setup, Agilio CX is compared with a traditional NIC for the number of VM instances, network throughput, packet latency and CPU

utilization.




Figure 2.

For performance comparison, a traditional NIC and Agilio CX (1x40G) is used on the eNodeB/MME and gateway (P/S-GW) servers. The User Equipment (UE) and Public Data Network (PDN) servers are used with a traditional NIC as they act as packet source and sink for mobile traffic.

Each of the VMs (VNF) run the following application software:

Tieto IP Stack (TIP) user space implementation

TIP is a portable, telecom grade RFC compliant stack, supporting the TCP/IP family of protocols, IPV4 and IPV6, in single or dual mode. TIP has a modular SW architecture, which allows easy interoperability with the standard Linux stack. TIP is used in a poll mode configuration on top of the DPDK.




Figure 1.



vEPC Application Software

The MME establishes UE sessions by sending requests over the S1-U interface to PDN. When a session is established via the MME emulator, all subsequent traffic is sent from the UE via the eNodeB and SGW/PGW to the PDN, and back through the same path. The PDN Gateway provides connectivity from the UE to external packet data networks by being the point of entry and exit of traffic, to and from the PDN.

vEPC Performance comparison: Agilio CX vs Traditional NIC

1. Number of VMs (VNFs)

Traditional NIC: A traditional NIC uses 3.3 cores per VM on the eNodeB and vEPC for the following:

- DPDK-thread inside the VM which is mapped to one of the VM threads on the host
- Vhost 1, relaying traffic between one port on the NIC to the hypervisor
- Vhost 2, relaying traffic between the other port on the NIC to the hypervisor
- Interrupt requests arriving from the NIC

Since 3.3 cores are used for the single VM, only 7 VMs can be created with 24 CPU cores.

Agilio CX: Since the Agilio CX is running the OVS offload Agilio software and using the PCI SR-IOV functionality, it does not have the overhead of the Vhost and interrupt processing. The VM implementation on the Dell server with Agilio CX takes only about 1.1 CPU cores. Therefore, on a dual-socket 24 core Dell R730 server up to 22 VMs can be implemented.

2. Throughput

Traditional NIC: For 128B packet sizes the traditional NIC can achieve 2.5Mpps (2.9Gb/s) with 7 VMs, the maximum possible on a traditional NIC, performs at about 350Kpps per VM.

Agilio CX: The Agilio CX SmartNIC can achieve performance of 13Mpps (15.5Gb/s for 128B packet) with up to 22 VMs on the same server, which is about 900 Kpps per VM (throughput flattens after 15 VMs due to other resource saturations in the server).

3. CPU Utilization

For the CPU load benchmarking, average load of 24 CPU cores and relative average load (load per Mpps) for 24 cores are measured for a traditional NIC and an Agilio CX.

Traditional NIC: For 128B packet sizes the relative peak load at highest rate achieved (2.5Mpps) is about 45% and the relative load is 18% per Mpps.

Agilio CX: For the comparison with a traditional NIC, at 2.5Mpps throughput graphs the load average is 12% and relative load average is 4% (CPU load per Mpps). At peak throughput of 13 Mpps, the average load is 52% and the relative load average is 4%.

4. Latency

The latency is measured as the end-to-end round trip time (RTT) latency between server and PDN server. Latency is measured by ICMP echo requests running in parallel to the traffic.

Traditional NIC: For the traditional NIC, 90% of 128B packets have below 1.5 milliseconds of latency.

Agilio CX: For 128B packet size data traffic, more than 90% of the 128B packets observe round trip latency time of less than 1ms, which is about 50% better than as observed in the case of the traditional NIC. In the case of other packet sizes, 256 bytes and larger, the latencies increase as they saturate the link bandwidth of 20Gb/s.

CONCLUSION

The performance improvements of the Agilio CX over the traditional NIC are highlighted in the table below.

PARAMETERS	TRADITIONAL NIC	AGILIO CX
Traffic Rate (Mpps) 128B	2.9 Gb/s	15 Gb/s
Peak Throughput at SGW-U instance (GTP) - 128B	2.5 Mpps	13 Mpps
Peak Throughput at PGW-U instance (GTP) - 128B	2.5 Mpps	13 Mpps
Number of UP VMs	7	22
Relative CPU Utilization (on the UP server) 128B	18%/Mpps	4%/Mpps
Latency UE ← → PDN, RTT, 128B	1.5 ms	1.0 ms

Note - The parameters are measured at 20G line rate (of total 40G available) for Traditional NIC and Agilio CX

The Netronome Agilio SmartNIC with Agilio software, which uses the SR-IOV and OVS offloads, provides more than 5X increase in the vEPC performance.

Agilio SmartNICs take advantage of OVS datapath offload with support for SR-IOV for direct communication from VMs to the Agilio CX hardware. This saves considerable CPU resources and enables 3X more VMs per server as compared to a traditional NIC. This translates to significantly lower data center server capital and operational expenditure.

