



White Paper

Delivering on the Promise of Server-Based Networking

Prepared by

Roz Roseboro
Senior Analyst, Heavy Reading
www.heavyreading.com

on behalf of

NETRONOME

www.netronome.com

August 2016

Introduction

With the advent of cloud technology, software-defined networking (SDN) and network functions virtualization (NFV), data center and telecom operators are transforming their data center architectures. They are moving away from purpose-built hardware to standardized commercial-off-the-shelf (COTS) hardware with applications running in virtual machines (VMs).

The Webscale Internet companies have proven that network virtualization can run on COTS servers along with other network functions – an approach called server-based networking. However, these companies also use proprietary acceleration techniques to get the requisite performance. Offloading the datapath to a Smart-NIC (network interface card) is proving to be a viable alternative, as it increases the number of cores that can support applications. At the same time, operators need to ensure that they can keep up with the pace of software innovation. In order to do so, they must use software that can be continuously updated, while remaining independent of the underlying hardware.

Netronome helps improve server-based networking by combining the performance of hardware with the innovation speed of software. Its networking software works in conjunction with intelligent server adapters (ISAs) to support compute and service node applications. These ISAs are fully programmable, so operators don't need to replace cards when they upgrade their software. With the software and hardware working in tandem, Netronome's platform aims to provide the performance and flexibility operators need in their points of presence (PoPs), enterprise data centers and telco/NFV cloud.

This white paper is structured as follows:

- **Section 2** discusses why cloud and telecom operators are transforming their data center architectures, along with some of the underlying technologies and some of the emerging challenges these operators face.
- **Section 3** explains why server-based networking will be critical in data center transformation, and how offloading the datapath can improve application performance. It also highlights the importance of continuous software updates.
- **Section 4** provides an overview of Netronome's software and hardware innovations, explaining how its solutions work in tandem to improve performance and agility.
- **Section 5** presents performance data for Netronome's solutions to show how it is able to deliver hardware performance at the speed of software.

Cloud & NFV Driving New Data Center Architectures

Data center and telecom operators are transforming their data centers to support SDN and cloud architectures, with telcos looking to support NFV as well. With this transition comes a move away from purpose-built proprietary hardware to commercial-off-the-shelf (COTS) servers. Applications running in VMs on those servers generate considerable East-West traffic, which leads to more networking (i.e., switching) functionality being deployed as well. Increasingly, with advancements in processor

technology, the switching intelligence is being supported on the servers themselves, rather than separate devices. These COTS servers were originally designed to support IT workloads, and not the more latency-sensitive telecom workloads. As a result, software-based and hardware-based acceleration techniques are being proposed to ensure COTS servers are able to provide comparable performance as dedicated platforms.

The advent of these new cloud-enabling architectures is leading to new challenges. For cloud/IaaS providers, sharing infrastructure among multiple tenants can lead to sub-optimal performance for individual tenants or applications. 4G/5G telco cloud service providers could see poor call or video quality when resources are not allocated properly. In both cases, operators are seeing massive growth in data center traffic and different traffic patterns, which makes problem resolution a challenge. Existing tools and processes take significant time, and are often expensive and/or inaccurate. While this is a concern in an IT environment, it becomes even more critical in a telecom environment, where not meeting service-level agreements (SLAs) can negatively impact operator revenues.

Server-Based Networking to Play a Critical Role

Overlays are an attractive option to data center operators because they reduce the need for frequent modifications to the physical network, while providing more agile deployment of networking resources. Overlay networking describes the creation of a virtual network that runs on top of a Layer 2 or Layer 3 physical network. Overlays are most often used to connect VMs to each other, but they are also used to connect VMs to gateways to the physical network including bare metal. They are created using various encapsulation techniques such as VXLAN.

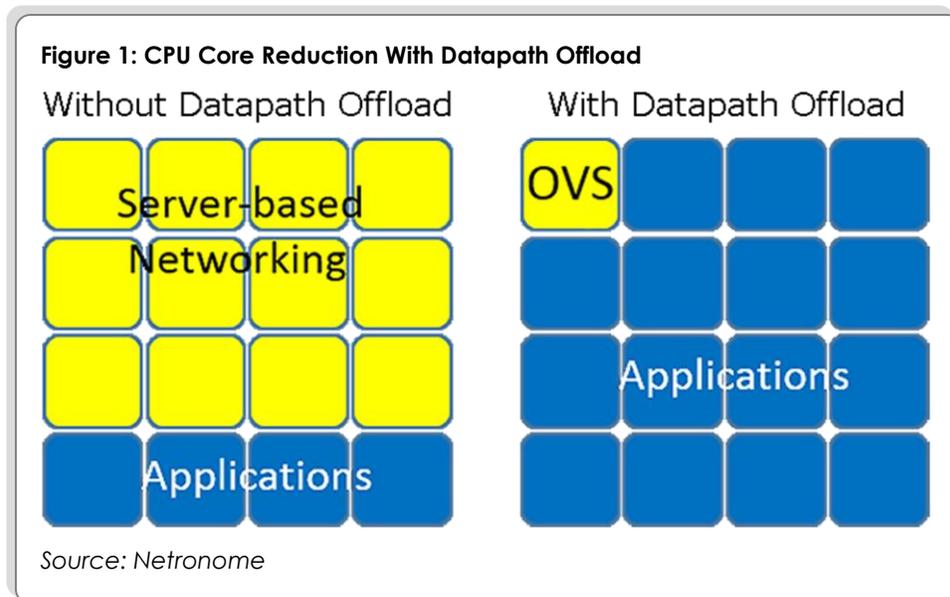
With server-based networking, network virtualization, as well as other functions such as security, load-balancing and telemetry are all run on a COTS server. Usually, these servers are managed using either an SDN controller or a cloud orchestration system. The Webscale Internet companies such as Microsoft, Google and Facebook are using server-based networking to support their public cloud services. The companies also have tremendous financial resources and specialized networking skills that most other companies lack. They also use proprietary silicon and leverage both hardware- and software-based acceleration techniques.

Datapath Offload Can Improve Application Performance

Server-based networking in x86 central processing units (CPUs) is an attractive option because of the flexibility it provides. The problem is that this can consume as many as 12 cores, which means four cores (or fewer) are left to support the applications running in VMs or containers. As a result, servers can only deliver 20 percent of the performance that they should, meaning servers are running at reduced efficiency.

One way to avoid this situation is to run the switching functionality on a network adapter or SmartNIC, rather than in the CPU. This frees up CPU cores so that they can be used to provide application processing, as shown in **Figure 1**.

Offloading the datapath in this manner will be key to allowing operators to achieve the performance they need to support SDN and NFV, and is especially important for speeds of 10 Gbit/s and higher.



Continuous Software Changes to Accommodate Networking Feature Adds

As discussed, the data center environment is becoming increasingly complex with the emergence of virtualization techniques. Routing, switching and firewalls are all being delivered virtually, and managed by new SDN controllers and cloud orchestration software.

Open source is driving many of the innovations in these technologies, but data center operators and vendors still need to find ways to differentiate their products and services. Historically, vendors write code on x86 servers using C, then port that code to programmable networking hardware. This process leads to very long development cycles, and sometimes unpredictable or poor performance. And because the code is dependent on the hardware, there is a significant risk of vendor lock-in.

Having hardware that can only support one function severely limits operators' ability to deploy new features as the underlying software is updated. This approach is proving to be a hindrance to operators competing with hyperscale and OTT players, which are operating under a DevOps model that provides them with greater agility – allowing them to react and deliver services more quickly.

Security Adds Complexity to Server-Based Networking

Security is another area of rising concern. Cloud architectures are more distributed than traditional architectures, so security, too, must be distributed. Policies must be defined at a VM level, which very quickly leads to a scale problem, with tens of thousands of rules taking up valuable processing resources on the server.

As operators move from VMs to containers, this challenge only increases, as the number of containers supported on a server will be substantially larger than the number of VMs. Sharing resources also makes it imperative to operate zero-trust security approaches.

Netronome Delivers Both Software & Hardware Innovations

Netronome's server networking solutions combine flow-processing-optimized silicon and a software architecture that runs on COTS hardware, and works in conjunction with standard, open-source network software. It continually updates its software to ensure that its customers can benefit from updates in community-driven projects such as Open vSwitch (OVS) without needing to change out their hardware. Netronome also provides tools to facilitate SDN application development.

Software Solutions

Netronome's Agilio software works in tandem with its ISAs to support compute and service node applications. Using open-source software such as OVS, Juniper's Contrail vRouter, Linux IP Tables and Connection Tracking, Netronome accelerates and scales network virtualization, security, load balancing and telemetry applications running on standard servers. Open vSwitch Database (OVSDB) interfaces and OVS command line interface (CLI) remain intact while the networking applications are offloaded, so its software can be provisioning using SDN controllers and cloud orchestration tools. Netronome's P4/C developer tools leverage the open source P4 compiler from the P4 Language Consortium.

Figure 2: Netronome Software Solutions

Product Name	Description
Agilio OVS	Supports multi-tenant policies required for cloud delivery. Offloads the entire vSwitch datapath
Agilio vRouter	Supports mobility services for SDN and NFV workloads using vRouter
Agilio OVS Firewall	Provides stateful security for VMs using OVN or Contrack
Agilio P4/C SDK	Provides the tools to allow developers to program SDN applications. It augments the open source compiler with extensions to enable optional C-based modules for more sophisticated functions.

Source: Netronome

Hardware Solutions

Netronome's Agilio CX family of ISA are fully programmable, meaning the software can be upgraded without replacing cards and development times can be dramatically reduced.

Agilio CX ISAs offload the virtual switch's datapath for networking functions to maximize the number of CPUs to be used for application processing. The Agilio CX family of ISAs fully and transparently offloads virtual switch datapath processing for networking functions such as overlays, security, load balancing and telemetry, enabling

compute servers used for server-based networking to deliver at their full potential. Agilio CX adapters are ideal for hardware-accelerated OpenStack networking, mobile core applications such as vEPC, and data center microsegmentation. These low-profile PCIe adapters fit in all COTS servers, including 1U configurations.

Achieve Hardware Performance at the Speed of Software

Netronome's Agilio platform provides the same low latency/high performance and CPU core savings as with SR-IOV, while also supporting a wide range of overlay tunneling protocols, including VLAN, VXLAN, GRE, MPLS and NSH.

Netronome's testing shows that its Agilio OVS solution can significantly improve throughput compared with OVS running in kernel or user space. Measuring throughput of OVS L2 Forwarding to VMs, using 100 Wildcard Flow Rules (match of MAC DA), 64,000 Exact Match Microflows and bursting with 10 packets/Microflow, Netronome achieved throughput of 28 Mpps at packet sizes of 64-512Kb, while the user space OVS achieved only 10 Mpps and kernel space OVS only 6 Mpps. It saw similar results with VXLAN overlays (where matching was done on VNI and MAC DA).*

While SR-IOV and DPDK on other NICs may support anti-spoofing, Netronome also supports stateless ACL security and stateful firewall, as well as VM migration and per-flow analytics. These features and capabilities extend the value of hardware-based efficiencies beyond NFV PoPs and enterprise data centers into the telco cloud and NFV environments.

Conclusion

Data center and telecom operators are transforming their data center architectures to support cloud, SDN and NFV, using COTS hardware as the underlying platform. They have seen that server-based networking can be used to support network functions, but are challenged to achieve the requisite performance.

By offloading the datapath to hardware, more cores can be devoted to application processing, and thus improve application performance. To keep up with the pace of change in networking features, software must be continuously upgraded without needing to change the underlying hardware.

Netronome's software works in conjunction with its fully-programmable Intelligent Server Adapters to bring the performance and flexibility that data center and telecom operators need to support their virtualization, automation and cloud services ambitions.

* Further information regarding this data and other testing results, along with the system configurations, can be found here: https://www.netronome.com/media/re-dactor_files/WP_OVS_Benchmarking.pdf