



OpenStack Networking: Where to Next?

WHAT IS STRIKING IS THE PERVASIVE USE OF OPEN VSWITCH (OVS), AND AMONG NEUTRON FEATURES, THE STRONG INTEREST IN SOFTWARE-BASED NETWORKING ON THE SERVER, OR SERVER-BASED NETWORKING.

INTRODUCTION

In April 2016, the OpenStack Foundation released a very enlightening [User Survey](#) which focused on networking features in OpenStack. The following key metrics (as percentage of respondents) from the survey came to the forefront:

- Usage: Full operational use, in production – 65%; On-premise, private cloud deployment – 65%; Neutron Networking – 90%
- Important emerging technologies: SDN/NFV – 52%; Containers – 70%; Bare Metal – 50%
- Workloads and frameworks: Private and public cloud products and services – 49%; NFV – 29%
- Preferred hypervisor: KVM – 93%; VMware – 8%
- Neutron drivers in Use: OVS – 80%; ML2 – 31%; Linux Bridge – 21%; Contrail – 14%; SR-IOV – 2%
- Neutron features actively used/interested in: Distributed across many features. Those that got 20% or higher: software load balancing, routing, VPN, accelerated vSwitch, QoS, software firewalling

What is striking is the pervasive use of Open vSwitch (OVS), and among Neutron features, the strong interest in software-based networking on the server, or server-based networking. Emerging technologies such as SDN/NFV and use of containers in bare metal server environments further justify the need for server-based networking. This trend, starkly visible from the survey report, is naturally in line with how the largest data centers in the world (such as Microsoft Azure) have opted for server-based networking as the most flexible and scalable form of networking.

With increasing bandwidth requirements and adoption of 10GbE and higher speed technologies, it is well understood today that accomplishing server-based networking tasks using general purpose CPUs is highly inefficient. Microsoft has indicated use of SmartNICs to scale their server-based networking infrastructure more efficiently. SmartNICs are used to offload and accelerate the server-based networking datapath. Ericsson and Netronome have demonstrated up to 6X lower TCO utilizing Netronome Agilio hardware-accelerated server-based networking incorporated in the OpenStack-based Ericsson's Cloud SDN platform.

To achieve such benefits, enhancements to the OpenStack networking plug-in architecture are needed. Let's start by exploring what is feasible today.



CURRENT STATE OF ACCELERATED OPENSTACK NETWORKING

At the OpenStack Summit 2015 in Vancouver, Mirantis presented the [state of hardware-accelerated networking](#). The below discussion is based on excerpts from that comprehensive presentation.

Today, SR-IOV plug-ins used with SR-IOV-capable server NICs are the only feasible form of networking hardware acceleration. When this is implemented in OpenStack Neutron using traditional NICs, bandwidth delivered to VMs is improved significantly, latency for VM to VM traffic is reduced, and CPU use for networking tasks is greatly reduced. However, the number of server-based networking features available at such high performance is severely limited. Figures 1a and 1b show the packet datapath from a network port on the server to a VM. With the datapath shown in Figure 1a, high bandwidth and low latency benefits are available for a very limited set of features as shown in Table 1, making the solution feasible for a very limited set of applications. When more features are needed, one needs to fall back to the datapath shown in Figure 1b, resulting in poor performance and high CPU usage. Note that SR-IOV adoption based on the recent survey is only 2% and this could be a reason.

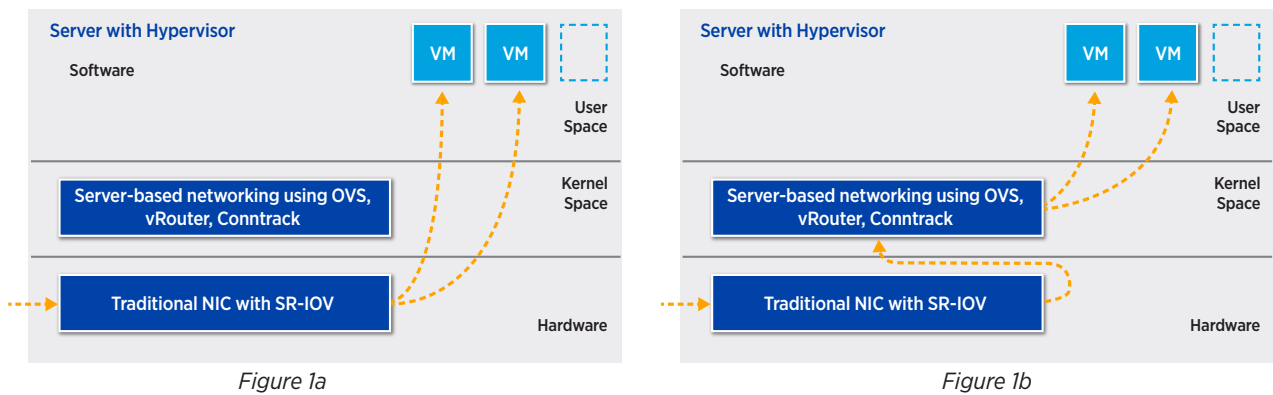


Figure 1a

Figure 1b

OpenStack Networking Plug-In Features	SR-IOV with Traditional NICs
VM Performance	High
Latency	Low
CPU Core Saving	High
Net Virtualization & Other	VLAN only
Anti-Spoofing	Yes
Load Balancing	No
ACL Security (stateless)	No
Firewall Security	No
VM Migration	No
Per Flow Analytics	No

■ Favorable
 ■ Unfavorable
 ■ Acceptable

Table 1

Some of the above challenges can be addressed using DPDK where the server-based networking datapath (as in OVS or Contrail vRouter) is maintained in software, but is moved to the user space to reduce user space to kernel space context switching overheads. This



approach is shown in Figure 2 below. This approach delivers additional server-based networking features (compared to SR-IOV) and improved performance (compared to a kernel-based networking datapath). This however comes at the cost of high latency and high CPU usage, and with use of larger number of flows in the networking datapath (as needed in security, load balancing and analytics applications) performance can be severely affected. As such, the DPDK approach has not found widespread adoption either.

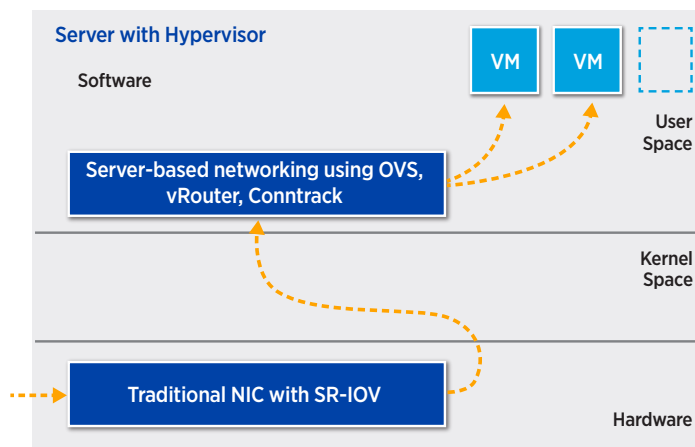


FIGURE 2

OpenStack Networking Plug-in Features	DPDK with other Vendor NICs
VM Performance	Med
Latency	High
CPU Core Saving	Low
Net Virtualization & Other	VLAN, VxLAN, GRE
Anti-Spoofing	Yes
Load Balancing	Yes
ACL Security (Stateless)	Yes
Firewall Security	No
VM Migration	Yes
Per Flow Analytics	No

■ Favorable
 ■ Unfavorable
 ■ Acceptable

This is the current state of accelerated OpenStack networking as was aptly summarized by the presentation at the OpenStack Summit 2016 in Vancouver. The April 2016 Survey sheds new light on what is needed from OpenStack networking. This leads us to the next question.



WHERE DO WE GO FROM HERE?

At the OpenStack Summit 2016 in Austin, Netronome demonstrated how OpenStack networking can be extended to deliver hardware-accelerated server-based networking. Netronome’s Agilio server networking platform will be utilized with OpenStack, OVS, Contrail vRouter and Linux Contrack technologies to showcase an extended set of features made possible at high performance (high bandwidth and low latency) and efficiency (low server CPU usage, freeing them up for more VMs). This is enabled by offloading the OVS, Contrail vRouter or Linux Contrack datapaths into the Agilio intelligent server adapter (SmartNIC) as shown in Figure 3. This approach, highlighted in Table 3, shows features needed in a broad array of workloads and deployment frameworks can be enabled at high performance and efficiency. This includes infrastructure as a service (IaaS), telco cloud SDN and NFV, as well as traditional IT and private/hybrid cloud workloads.

NETRONOME’S AGILIO SERVER NETWORKING PLATFORM WILL BE UTILIZED WITH OPENSTACK, OVS, CONTRAIL VROUTER AND LINUX CONTRACK TECHNOLOGIES TO SHOWCASE AN EXTENDED SET OF FEATURES MADE POSSIBLE AT HIGH PERFORMANCE AND EFFICIENCY.

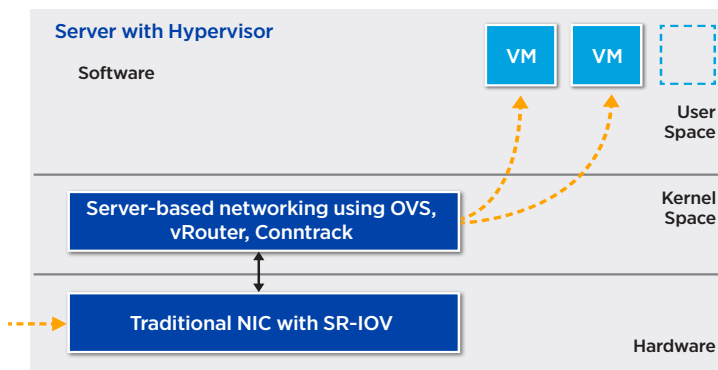


Figure 3

OpenStack Networking Plug-In Features	SR-IOV With Other Vendor NICs*	DPDK With Other Vendor NICs*	Enhancements Based On Netronome Agilio Platform
VM Performance	Med	Med	High
Latency	High	High	Low
CPU Core Saving	Low	Low	High
Net Virtualization & Other	VLAN only	VLAN, VxLAN, GRE	VLAN, VxLAN, GRE, MPLS, NSH
Anti-Spoofing	Yes	Yes	Yes
Load Balancing	Yes	Yes	Yes
ACL Security (Stateless)	Yes	Yes	Yes
Firewall Security	No	No	Yes
VM Migration	Yes	Yes	Yes
Per Flow Analytics	No	No	Yes
Markets	NFV POP Clusters	On-premise NFV	POP, IaaS & Enterprise DC, NFV

■ Favorable
 ■ Unfavorable
 ■ Acceptable



This model of delivering server-based networking using SmartNIC hardware such as Netronome’s Agilio platform is also ideally suited for bare metal server-based container deployments of the future. Features such as networking virtualization, service chaining, load balancing, security and analytics can be implemented and provisioned from outside the domain of the operating system running on the server. In this case, control plane orchestration using an SDN Controller or OpenStack is implemented directly with the SmartNIC. This concept is shown in Figure 4 below.

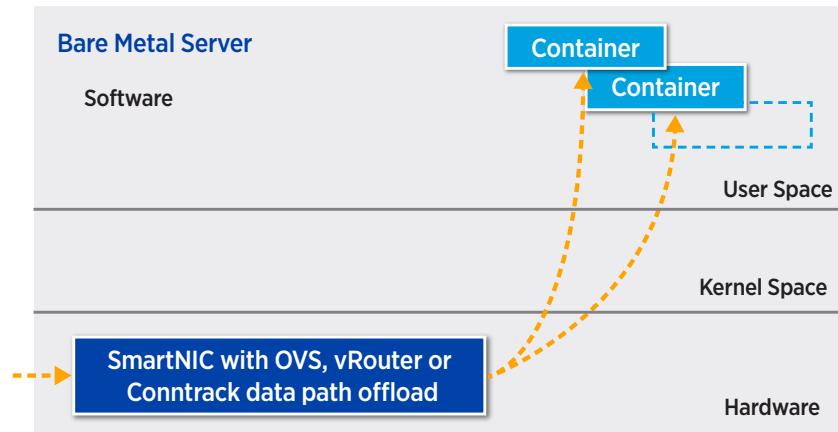


Figure 4

CONCLUSION

For all of this to be a reality, the OpenStack networking plug-in specification will have to be enhanced beyond the SR-IOV-based capabilities that exist today. Netronome is taking a leadership role in the development of these enhancements with industry leaders such as Mirantis, Ericsson, Juniper Networks and others. A draft open specification covering these enhancements is expected in Q3 of 2016 and this will be contributed to the OpenStack community for further feedback and eventual acceptance for industry-wide adoption.