

Agilio vRouter Software

ACCELERATING SERVER-BASED NETWORKING FOR JUNIPER CONTRAIL VROUTER



Network automation and configuration are not disrupted with Agilio, making the user experience identical to software-only deployments – the only difference is performance. Users will experience a 6X gain in performance with Agilio-accelerated vRouter platforms over software-only vRouter platforms.

Overview and Benefits of Contrail Networking

Contrail is a scale-out, standards-based virtual networking solution that seamlessly integrates with physical routers and switches to eliminate the challenges of private and public cloud networking. The solution primarily targets cloud networking in the form of private clouds, Infrastructure as a Service (IaaS), and virtual private clouds (VPCs) as well as Network Function Virtualization (NFV) in service provider networks.

A key networking use case for Contrail is overlay networking, which addresses the physical networking challenges associated with rapid configuration, provisioning, automation, and overall scalability in today's data center networks. Contrail includes a logically centralized but physically distributed controller, and dataplane vRouter that resides on each compute node within the data center. The Contrail vRouter is responsible for extending data center networking to the server and managing overlay networking. Virtual overlay networks are created on top of the physical underlay network. In Contrail, the overlay tunnels are created with MPLS over GRE, MPLS over UDP, or VXLAN encapsulations.

Using these encapsulations, MPLS over GRE/UDP specifically, brings several advantages. First, MPLS (over GRE) is widely supported by existing routers from all major vendors. Second, these encapsulations draw many similarities to MPLS L3VPNs and EVPNs, which are mature and known to scale. Third, because the MPLS technology is prevalent in production networks today, Contrail will seamlessly interoperate with existing network equipment without the need for software gateways for network translation and mapping.

The Challenges of Server-Based Networking

There are new challenges that Contrail imposes, however, that threaten data center performance. Contrail vRouter significantly increases the workload associated with server-based networking. In software-only vRouter deployments x86 cores are responsible for executing overlay networking tasks, having a negative effect on performance. The penalties associated with software-only networking on x86 are as follows:

- 86 CPU cores are not optimized for server-based networking, which result in low packet rates and low network bandwidth
- Valuable resources are lost due to server-based networking – x86 cores are allocated to networking tasks rather than hosting VMs
- I/O bottlenecks are present due to expensive memory copies associated with hypervisor networking

The typical Contrail vRouter deployment includes four (4) x86 cores, with each core supporting on average 1Mpps of packet forwarding (workload dependent). Additional cores can be assigned, however there must be a balance between cores dedicated to networking and cores dedicated to virtual machines (VMs). This common configuration allows for 4Mpps of vRouter forwarding and traffic availability to VMs. When counting for bi-directionality, this translates to 2Mpps of receive traffic and 2Mpps of transmit traffic for VMs. In many cases, VNFs within VMs can process roughly 1Mpps per instance (or core), with as many as 20 or more VNFs deployed per server. The mismatch in packet rates show that VMs and VNFs are severely undersubscribed with software-only networking.

A very important benefit of the Contrail solution is enabling the ability for data centers to seamlessly scale out by pushing networking tasks such as switching, routing, load balancing, and firewalling to compute nodes. However this benefit causes a server scale-up problem. Because x86 CPU cores are not optimized for server-based networking, operators lose the ability to scale-up the performance of single servers. This has damaging effects on deployment and cost efficiencies across the data center.

Agilio vRouter Solution resolves the Scale-Up Problem

Netronome has provided the ability to accelerate Contrail vRouter and preserve valuable CPU cycles through offload to Netronome's Agilio-CX intelligent server adapters. The Contrail vRouter data path and its associated lookup tables are mirrored to the intelligent server adapter through the use of the Agilio vRouter software. Together, the Agilio hardware and software solution improve vRouter performance for VNFs and VMs significantly. Overlay networking performance is increased to 27Mpps with Agilio, compared to 4Mpps for software-only vRouter on x86. (Figure 1.)

For VMs and VNFs, native hardware I/O performance is available through the use of SR-IOV. This is a unique offering when considering the use of vRouter on x86 with traditional NICs. Because traditional NICs do not have the ability to offload the complex workloads of the vRouter data path, direct connectivity to VMs using SR-IOV is not possible when network tunneling is required. As a result, software-only implementations must rely on hypervisor-based networking, which require expensive memory and CPU operations that limit I/O performance. Agilio vRouter software and adapters eliminate these bottlenecks by connecting SR-IOV virtual functions (VFs) directly to the offloaded vRouter data path.

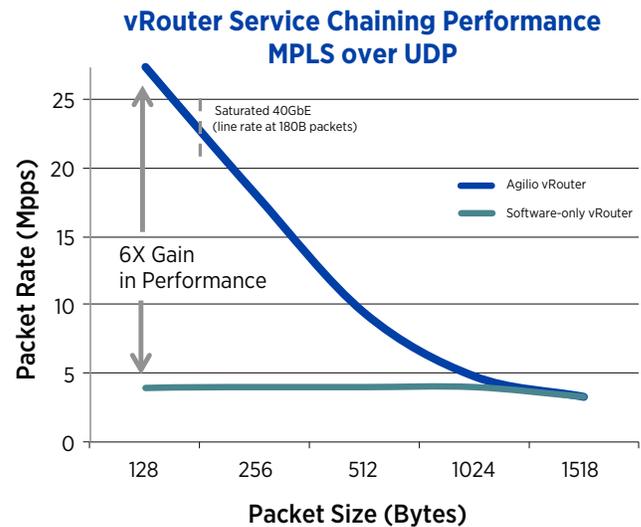


Figure 1. Packet Rates with Agilio vRouter and Software-only vRouter

Seamless Integration with Contrail

The Agilio-CX adapter and software will transparently receive configuration directly from the vRouter data path hosted in the kernel. As a result, the vRouter functionality, configuration, and control channels are completely preserved. This transparent offload allows the Agilio solution to integrate with the Contrail controller and vRouter subsystems seamlessly. (Figure 2.)

The common tunneling configurations that Agilio vRouter supports for L2 and L3 overlay networking include MPLS over GRE, MPLS over UDP, and VXLAN. This is accomplished by offloading and accelerating lookups and actions for several Contrail vRouter tables:

- Interface Tables
- Next Hop Tables
- Ingress Label Manager (ILM) Tables
- IPv4 FIB
- IPv6 FIB
- L2 Forwarding Tables
- Flow Tables

In the case of encapsulation, a VM or VNF would transmit a packet on a VF that is connected to the Agilio adapter. The Agilio vRouter data path receives the packet and executes the required lookups. In the case of an L3 overlay and network tunneling with MPLS over UDP (see figure 3), the inner

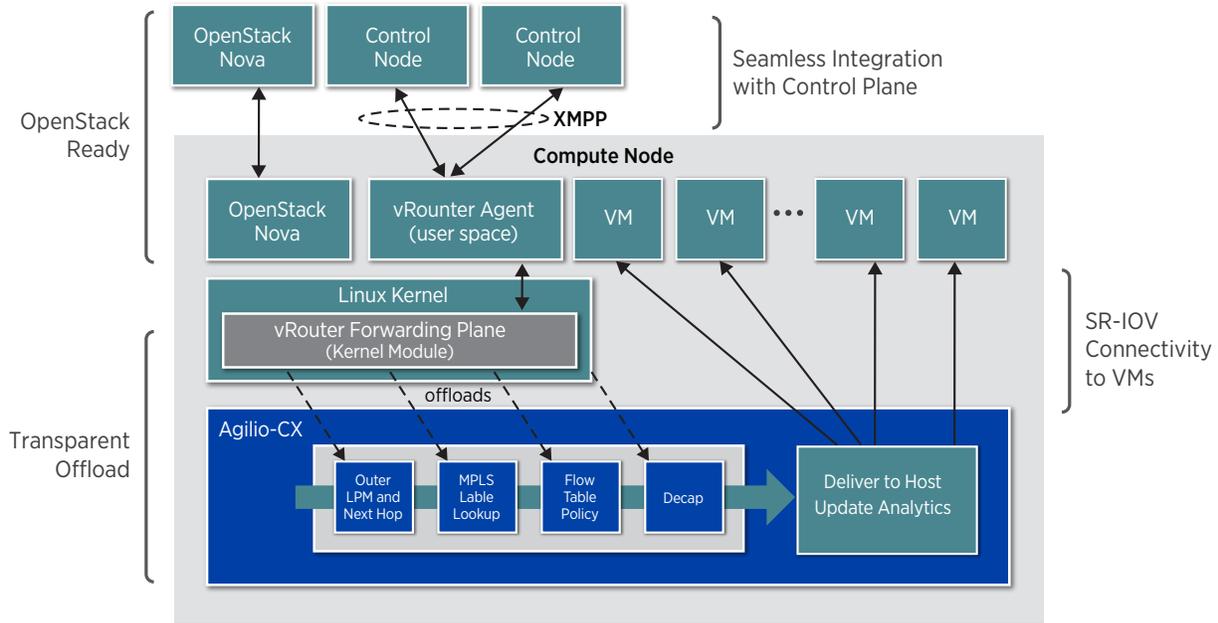


Figure 2 Compute Node Architecture and Control Plane with Contrail vRouter and Agilio Acceleration

Ethernet header is stripped and the MPLS over UDP headers are added for the tunnel. Finally, the underlay network headers (Ethernet and IP) are added based on MPLS label lookups and the final frame is transmitted on the network.

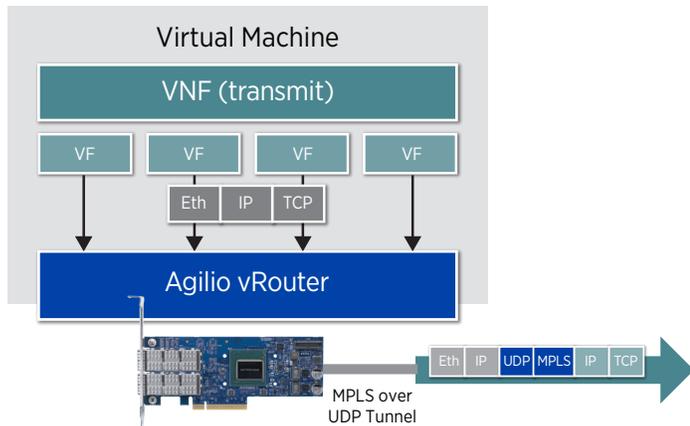


Figure 3. Encapsulation for Overlay Networking

Summary and Benefits

From a performance perspective, the Agilio vRouter solution allows operators to more fully utilize compute platforms through offload due to:

1. Increased PPS capacity of vRouter data path (27Mpps with Agilio) allows applications to reach their full processing capacity
2. Reclaimed CPU cores from vRouter offload allows more VMs and VNFs to be deployed per server
3. Native networking I/O performance to VMs and VNFs through the use of SR-IOV

Lastly, the functional benefit to Contrail users is that they can take advantage of the Agilio performance gains while retaining their pre-existing Contrail architecture. The control plane (XMPP) configuration for the vRouter data plane remains the same, and Open Stack management is unchanged. Network automation and configuration are not disrupted with Agilio, making the user experience identical to software-only deployments – the only difference is performance. Users will experience a 6X gain in performance with Agilio-accelerated vRouter platforms over software-only vRouter platforms.



NETRONOME

Netronome Systems, Inc.
2903 Bunker Hill Lane, Suite 150
Santa Clara, CA 95054
Tel: 408.496.0022 | Fax: 408.586.0002

www.netronome.com

©2016 Netronome. All rights reserved. Netronome, the Netronome logo, and Agilio are trademarks or registered trademarks of Netronome Systems, Inc. All other trademarks mentioned are registered trademarks or trademarks of their respective owners in the United States and other countries.
SB-vRouter-02/16