



OpenStack Support in Agilio®

AGILIO SMARTNICs
WITH AGILIO SOFTWARE
PROVIDE COMPLETE
INTEGRATION WITH
OPENSTACK. THE KILO
AND LIBERTY RELEASES
OF OPENSTACK ARE
SUPPORTED WITH AGILIO
PLATFORM.

OpenStack is an open source cloud-computing platform used by many telcos and service providers in private, public and hybrid clouds for providing automated and scalable compute, storage and networking resources. It is a collection of open source software modules running as a service on an integrated platform. The compute (virtual machines), networking and storage resources can be managed through command line tools, RESTful APIs or the web-based GUI called Horizon.

Agilio SmartNICs with Agilio Software provide complete integration with OpenStack. The Kilo and Liberty releases of OpenStack are supported in the Agilio platform. Users can follow Agilio OpenStack installation guide to build/install an OpenStack environment with multiple Netronome Agilio SmartNIC-powered compute nodes. On the Agilio SmartNICs users can perform the following operations with OpenStack:

- VM provisioning using SR-IOV
- VM startup using SR-IOV
- VM shutdown using SR-IOV
- VM migration (Supported in upcoming Agilio Software release)
- Configure the OVS rules using SR-IOV

For the managed installation of OpenStack nodes and services, Netronome has partnered with Mirantis to provide the GUI-based, managed automated installation.

OpenStack support in Agilio is implemented by modifying the software components interfacing with the OpenStack services; most of these modifications are done in the libvert driver for compute nodes used for the management of the VMs. Libvert interacts with the OpenStack Nova compute service for VM provisioning using the OpenStack Horizon GUI.

It is important to understand the libvert utility before explaining the OpenStack support for the Agilio SmartNIC as most of the OpenStack compatibility changes are implemented using the libvert software modification.

Libvert: libvert is collection of open source software that provides the APIs to manage the virtual machines. There are three parts of the libvert software:

- libvertd – daemon
- libvert – API calls
- libvert command line utility – virsh



THE AGILIO SOFTWARE IS MODIFIED TO PROVIDE MAC ADDRESS AND VLAN TAG FOR THE VMS TO BE USED BY THE LIBVERT UTILITY..

The libvert utility is supported on all common hypervisors, which include KVM, Xen, ESXi, qemu, libvert eliminates the need of using the hypervisor specific tools as it exposes all the useful hypervisor features. VM configurations are stored in a XML files (depending on the hypervisor) using the libvert CLI, XML files can be edited to modify the configurations and get validated by the libvert daemon. VM migration is also supported with libvert CLI when the libvert daemon is running on source and destination hosts.

The intelligent PCI pass-through functionality in the libvert retrieves the current MAC and VLAN tag for an SR-IOV VF from the VF configuration BAR and sets it as libvert configuration item. The Agilio Software is modified to provide MAC address and VLAN tag for the VMs to be used by the libvert utility.

The following requirements in the Agilio Software will be supported for the OpenStack integration:

1. Agilio Software will have a mechanism to set the MAC of the VF and fallback device when VM boots.
2. Libvert is patched to:
 - a. Pick an unused VF from predefined pool
 - b. Set the MAC of the VF
 - c. Attach the VF to a VM
 - d. Rename the fallback device to specified name
 - e. Add the fallback device to configured OVS bridge with specified options
6. Libvert will have a mechanism to define the VF pools to be used.
7. OpenStack services (Nova) will be modified to make aware of the VFs to use them instead of the standard Virtio/vHost.

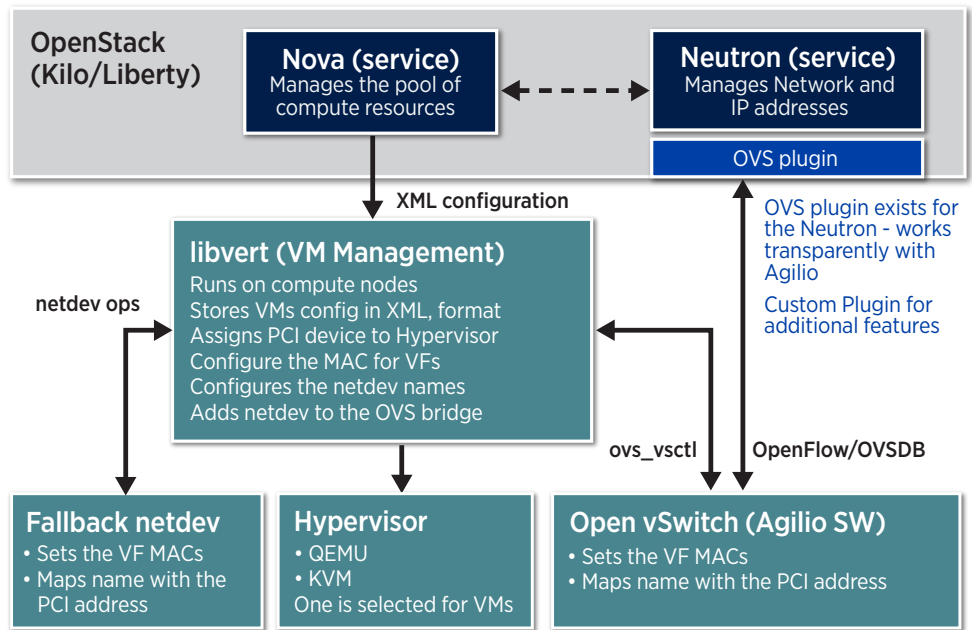


FIGURE 1 - OpenStack components for Agilio Software



The following blocks in Agilio are updated to support OpenStack:

- Agilio PCIe driver sets the VF's MAC address via the PF on next file level restore (FLR)
- Fallback netdev sets the future MAC and keeps a mapping of VF PCI addresses and their respective netdev names
- Libvert performs the following tasks by interfacing with Fallback netdev, Hypervisor and Agilio OVS software:
 - Assign PCI device to the hypervisor
 - Configures the future MAC address and netdev names
 - Adds the netdevs to the OVS bridge
- Nova compute service is part of the OpenStack software; it uses libvert to start VMs and also interfaces with the Neutron (OpenStack network service) to provide and verify the network configuration. It is updated with the changes made in the libvert module.
- Interfaces between the blocks are modified accordingly to support the functionality described previously. The modified interfaces include netdev-ops and libvert XML-config.
- The Neutron ML2 plugin for the OVS agent is already available and works transparently without the need of any modifications.
- SR-IOV ML2 plugin for the Netronome Agilio CX will be modified to use SR-IOV with the OVS in upcoming Agilio Software release.
- NAT and Security groups for the VMs will be supported with OVS Conntrak. The current OpenStack OVS ML2 plugin does not support the OVS conntrak. It will be supported in the future releases of the Agilio, which might be included with the OVN support in Agilio Software.

The Agilio platform provides the complete integration with OpenStack by modifying the open source components and Netronome Agilio Software. The changes made in open source software components (Nova, Neutron and libvert) will be upstreamed. In future releases, Agilio Software will include more VM configuration features such as NAT, security groups and VM migration.

Mirantis OpenStack (MOS) with Agilio SmartNIC

Mirantis provides an OpenStack package, which is updated with each stable release of OpenStack, which include all the services included with the release. Mirantis also provides managed OpenStack life cycle management software Fuel, which can deploy and manage multiple OpenStack environments. Mirantis OpenStack distribution supports wide range of operating systems and hypervisors.

To deploy the Mirantis OpenStack in a managed environment, one needs to use Mirantis Fuel software. Fuel acts as a single control plane for OpenStack deployment. Fuel plugins add to Fuel the ability to deploy OpenStack with partner software features and drivers, eliminating the need for manual configuration and patching. A number of Mirantis partners Fuel plugins already exist today. Netronome provides the Fuel plugin for the Agilio SmartNIC, which allows Mirantis OpenStack installation and orchestration using the Fuel UI and command line inter-



face. The following services can be configured on Agilio SmartNIC using the Mirantis OpenStack platform:

- Contrail vRouter
- Linux Firewall
- Open vSwitch (OVS)

Since all of the above services are offloaded to the Agilio SmartNIC transparently from the user orchestrating them using the OpenStack platform this device can deliver up to 5X higher networking throughput and 6X lower total cost of ownership (TCO).

vRouter OpenStack Integration with Agilio

Juniper's Contrail vRouter provides the virtual network that has OpenStack plugin. Netronome is providing the Contrail vRouter functionality offload on the network flow processor (NFP). The purpose of the NFP vRouter integration with OpenStack is to use the hardware-accelerated vRouter functionality through the OpenStack interface without knowing the details of the vRouter or the NFP internals.

The functional goal is to allow users to provision the VMs attached to NFP-accelerated interfaces, using a process that appears to be exactly identical to the process for provisioning VMs with non-accelerated interfaces. The accelerated NFP vRouter can work as a drop-in replacement to an unaccelerated vRouter. The OpenStack integration works equally well from the OpenStack dashboard ("Horizon GUI"), OpenStack command line (nova boot), and, if applicable, from the Contrail Web UI. The initial release of the NFP-accelerated vRouter will target OpenContrail version 3.0 and OpenStack Kilo.

The Contrail vRouter plugin for OpenStack is supported by the Agilio SmartNIC with modifying and adding the following OpenStack components:

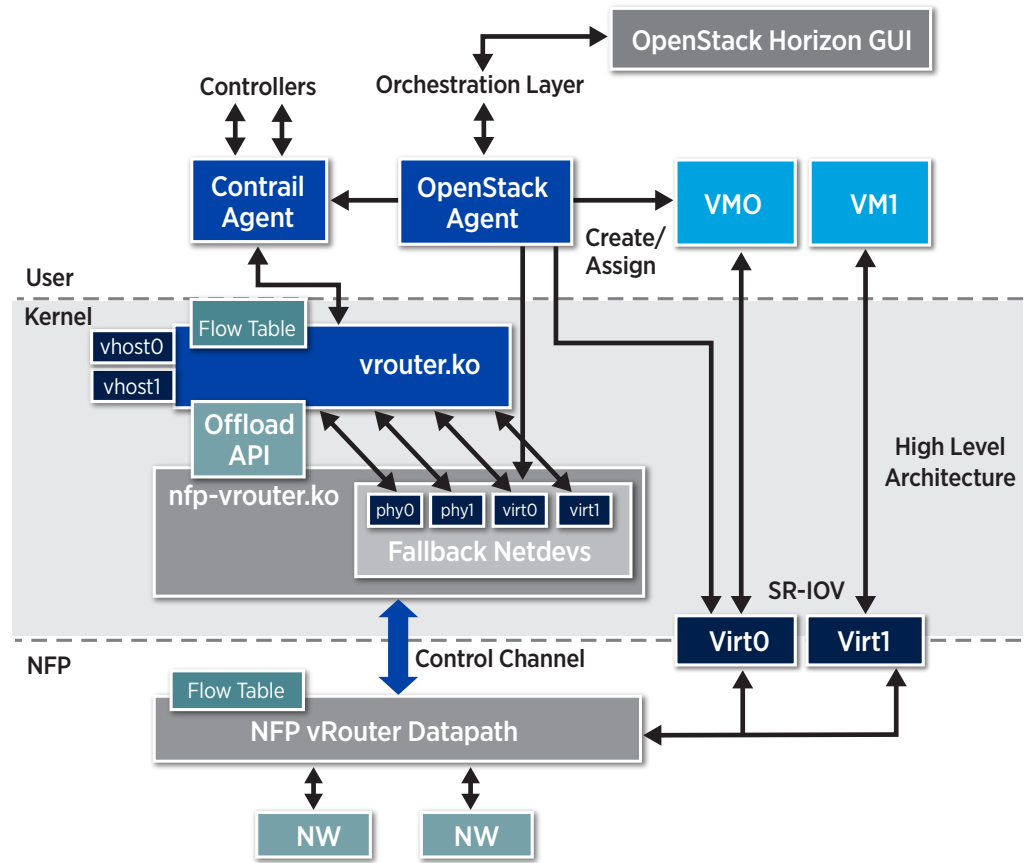
- Libvert daemon, running on the compute node. The libvert modifications will interact with the metadata from the contrail plugin.
- Nova compute modifications, which interfaces with the Libvertd.
- Contrail-vRouter-agent OpenStack agent for Contrail vRouter already exists, this plugin should work transparently while vRouter functionality is offloaded/accelerated on the Agilio SmartNIC.
- Accelerated vRouter kernel module developed by Netronome for vRouter offloaded functionality.

Netronome vRouter acceleration driver and modification of the OpenStack components will enable the user to configure the vRouter functionality through the OpenStack Horizon GUI, Few of the main vRouter features include:

- L3 forwarding
- Underlay and overlay networking (IPv4 and IPv6 overlay with IPv4 underlay)
- Tunnels (MPLS over GRE and MPLS over UDP)
- IP checksum offload



The figure below shows the vRouter integration with the Agilio platform (Agilio SmartNIC and Agilio Software) and their integration with the OpenStack.



As described above the Netronome Agilio SmartNIC platform provides complete integration with OpenStack orchestration where data center operators can orchestrate OVS and Contrail vRouter features using the OpenStack GUI, CLI or API options, which will provide easy configuration of switching and routing between the VMs.