

Enabling NFV With High Performance NFVI

Using generic servers, virtualization and high speed Ethernet technologies, **Network Functions Virtualization (NFV)** enables the transformation of telecom architecture. Instead of relying on dedicated equipment for each function, Telecom Operators can now design faster, more flexible architectures interconnecting Virtual Network Function (VNF) appliances running on generic servers.

Network Functions Virtualization Infrastructure (NFVI) is a key component of the NFV architecture that describes the hardware and software components on which virtual networks are built. NFVI creates a virtualization layer that sits right above the hardware and abstracts the hardware resources, so they can be logically partitioned and provided to the VNF to perform their functions.

NFVI is critical to deliver the required bandwidth to VNFs and Telecom Operators have specific functional and performance requirements for the NFVI:

- High performance switching aggregated bandwidth for VNFs
- Hardware independent VNF network attachments for seamless network hardware upgrades and VNF migration
- Communication between VNFs and NFVI using standardized interfaces that avoid to modify the VNF and provide VNF vendor independence

- High capacity, low latency inter-VNF communications to future proof service chaining capabilities
- Rich networking feature set to provide advanced services to VNFs beyond switching capabilities

To meet all these requirements, 6WIND Virtual Accelerator is designed around an open, high performance software switch to provide a high performance and full-featured networking platform for the NFVI.

Based on a fast path architecture, 6WIND Virtual Accelerator transparently accelerates the Linux selected software switch to provide extreme bandwidth and low-latency to VNFs. As an example, 6WIND's Open vSwitch (OVS) acceleration provides 20 Gbps of switching capacity per core that scales linearly with the number of cores without any modification to OVS or its management.

Using DPDK to interface with Ethernet Network Interface Cards (NIC), 6WIND Virtual Accelerator isolates VNFs from real hardware, thus eliminating driver dependency in VNF operating systems. This isolation also allows network hardware changes in either connection speed, supplier or function (for example a smart NIC performing DDoS mitigation) without VNF changes. It also allows for live and scheduled migration of VNFs from one platform to the other without any hardware NIC adaptations.

On the application side, 6WIND Virtual Accelerator interfaces with the VNFs using standard Virtio. The VNF can run any application, any Operating System without any changes. It allows Telecom Operators to easily integrate VNFs from different vendors with various guest operating systems.

Inter-VNF communications with high capacity and low latency are key differences between traditional Cloud and Telecom Operator environments. While east-west communication exists in Private and Hybrid Clouds, the performance characteristics are dramatically different and thus require a carrier-grade solution. 6WIND Virtual Accelerator supports sub-microsecond low latency with several hundreds of Gbps throughput of inter-VNF traffic. It does this without consuming external ports, impacting available bandwidth on the existing ports, or incurring PCI Express contention and latency.

Beyond accelerating Linux Bridge and Open vSwitch, 6WIND Virtual Accelerator also provides enhanced networking services such as overlay, L3 forwarding, virtual routing,

multi-tenancy, IPsec and more to extend service chaining capabilities between VNFs.

Transparent Deployment

6WIND Virtual Accelerator software deploys transparently into the hypervisor domain used by the NFVI based either on open source or commercial distributions. It supports OpenStack and SDN Controllers without any changes to the existing management and orchestration environment.

6WIND Virtual Accelerator is installed using standard Linux package managers, rpm or deb. In OpenStack environments, 6WIND Virtual Accelerator can be also installed using platforms such as Fuel (plug-in) or Juju (Charms).

6WIND Virtual Accelerator includes a ready-to-use “zero conf” default configuration option. Optimizations can be done using configuration tools. Once configured, 6WIND Virtual Accelerator is started as a service of the Linux distribution.


6WIND Virtual Accelerator for NFVI

Accelerates Open vSwitch (OVS) and Linux bridging to deliver an aggregated bandwidth of 200 Gbps to VNFs.

The platform that sustains this level of performance is a 10 core Intel processor server equipped with 40G NICs and running a standard Linux distribution and DPDK (Data Plane Development Kit).

Flows are configured by a standard and unmodified OpenFlow controller or OpenStack environment.

Beyond accelerating switching, 6WIND Virtual Accelerator also provides enhanced networking services such as overlay, L3 forwarding, virtual routing, multi-tenancy, IPsec and more.



6WIND VSR