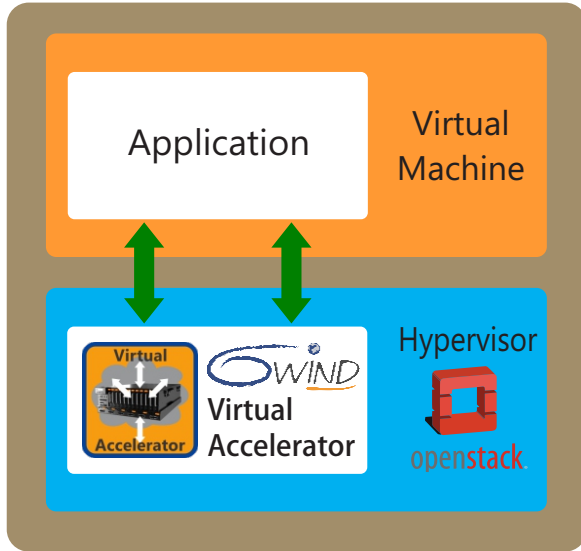




Virtual Accelerator™



Hypervisor Acceleration Software



Forwarding Performance¹

- 20 Gbps per core
- Performance scales linearly with number of cores

Highlights

- Increase VM and VNF performance without SR-IOV
- Multiply VM density on Commercial-Off-The-Shelf (COTS) servers
- 1/10/40/100GE support
- Transparent Operation: No change to existing VNFs or VMs or management tools (Linux, KVM, OpenStack, SDN)
- Zero hardware or software dependencies

Product Overview

6WIND Virtual Accelerator is part of the 6WIND vRouter family and provides packet processing acceleration for virtual network infrastructures used in Network Function Virtualization (NFV), Data Center Virtualization or Network Appliance Virtualization use cases. It is ready-to-use software installed on virtualization platforms.

6WIND Virtual Accelerator runs inside the hypervisor domain and removes Linux performance bottlenecks by offloading virtual switching from the Linux networking stack.

Virtualization on COTS Servers

6WIND Virtual Accelerator provides efficient switching to increase virtual infrastructure performance. Testing in live OpenStack environments shows popular Virtual Network Functions (VNFs) can achieve performance boosts of up to 5x and 20 Gbps throughput without SR-IOV or any changes to the application.

By offloading network packet processing to dedicated CPU cores, 6WIND Virtual Accelerator maximizes the number of Virtual Machines (VMs) without affecting individual networking performance requirements. The increased VM density can be applied to additional VNFs such as load balancers, firewalls, etc. to generate new revenue streams on existing servers.

Product Description and Architecture

6WIND Virtual Accelerator boosts virtual networking performance by running its data plane packet processing software on dedicated cores. The data plane software comprises DPDK, a set of poll mode drivers that provide efficient data I/O processing. In addition, 6WIND technology has native multi-queue support to distribute the load across multiple CPU cores. This combination allows 6WIND Virtual Accelerator performance to scale linearly with the number of cores, while Linux OVS is unable to scale.

6WIND Virtual Accelerator does not require any specific configuration and takes advantage of the standard Linux control plane (OVS, routing, etc.) to configure its forwarding tables. It uses a continuous synchronization mechanism that is completely transparent to Linux, so the control plane applications operate without any changes. 6WIND Virtual Accelerator transparently adds statistics regarding packets processed to Linux statistics, so no changes are required to monitoring tools.



NFV Of The Year
#SPEEDMATTERS
6WIND.com

¹Test Platform: Intel Xeon E5-2690 v3 @ 2.7GHz; test performed by 6WIND in a controlled environment

6WIND Virtual Accelerator Switch Capacity vs. Linux OVS; Intel^R Xeon^R E5-2690 v3 @ 2.60 Ghz (Mpps)

# CPU Cores	1	2	4	8	16	20
6WIND VA	7.8	12.6	22.2	41.4	79.8	99
Linux OVS	1.4	1.8	2.5	3.9	6.7	8.1

Key Features

- Mature solution for high performance I/O leveraging DPDK, with multi-vendor NIC support
- Accelerated virtual switching (Open vSwitch and Linux bridge) and networking (Overlays, Filtering/NAT, IP Forwarding, IPsec) using the fastest packet processing algorithms
- Preserves standard Linux commands (tcpdump, iptables, IPsec/XFRM, etc.) for ease of deployment and operation
- Supports standard network management tools
- Supports existing workloads through standard Virtio drivers
- Scalable software-based switching enables Service Chaining
- Performance without SR-IOV
- Advanced monitoring through streaming telemetry with time series data base and graphical analytics

System Requirements

CPU

- Intel x86 Xeon and Atom

Core count

- Minimum one and fully customizable

NICs

- Intel 1G 82575, 82576, 82580, I210, I211, I350, I354
- Intel 10G/40G 82598, 82599, X520, X540, XL710
- Mellanox 10G/25G/40G/50G/100G CX3, CX4, CX5
- Broadcom NetExtreme E-Series

Supported Linux distributions

- Red Hat Enterprise Linux
- Ubuntu
- CentOS

Supported OpenStack distributions

- Ubuntu Cloud and RDO

Performance

- Switching Capacity (OpenStack)
 - Compute Node: 20 Gbps per core
 - Network Node: 40 Gbps per core
- Forwarding: 12 Mpps per core
- VLAN: 10 Mpps per core
- Filtering: 6 Mpps per core

Features

Virtual Switching

- Open vSwitch (OVS)
- Linux Bridge
- Virtio Host poll mode driver (PMD) allows communication with any workload with Virtio drivers (can be DPDK, Linux or other OS-based)

L2 and Encapsulations

- GRE
- VLAN (802.1Q, QinQ)
- VXLAN
- LAG (802.3ad, LACP)
- Ethernet Bridge

Offloads (for Virtio guest)

- Checksum offload (IP and TCP/UDP)
- LRO (based on GRO)
- TSO (based on GSO)

These offloads work on inner headers within tunnels (VLAN, VXLAN, GRE, IPinIP). Virtual Accelerator leverages hardware offloads when supported by the NICs.

IP Networking

- IPv4 and IPv6
- VRF
- IPv4 and IPv6 Tunneling
- NAT

- ARP Proxy²

Management/Monitoring

- Standard Linux networking tools: ovs-*, iproute2, iptables, ping, tcpdump, traceroute, telnet
- Syslog
- KPIs/Telemetry (YANG-based)
- SNMP

Security

- ACLs (stateless and stateful)
- uRPF
- CP Protection (including rte flow HW support²)

IP Services

- DHCP Server/Client/Relay
- DNS Client/Proxy
- NTP

QoS

- Rate limiting per interface
- Rate limiting per VRF
- Class-based QoS
 - Classification: ToS/IP/DSCP/CoS
 - Shaping and policing
 - Scheduling: PQ, PB-DWRR

Pricing

- Pricing is based on switching capacity: 20G, 50G, 80G, 200G
- Switching capacity = Traffic (uplink and downlink) for each physical NIC plus inter-VM traffic (traffic chaining)

²Near-term Roadmap

6WIND Speed Certified Applications



6WIND Speed Certified Infrastructure And Platforms

