WHITE PAPER

Hardware Versus Software Switching

Any infrastructure will become obsolete over time, and need to evolve to adapt to changing usage needs. Few industries are impacted as directly by the need to evolve as the networking industry.

Telecom operators are pressing technology providers to provide them with the means to virtualize network nodes that accomplish certain network functions on top of a generic, open virtualized infrastructure: the Network Function Virtualization Infrastructure, or NFVI.

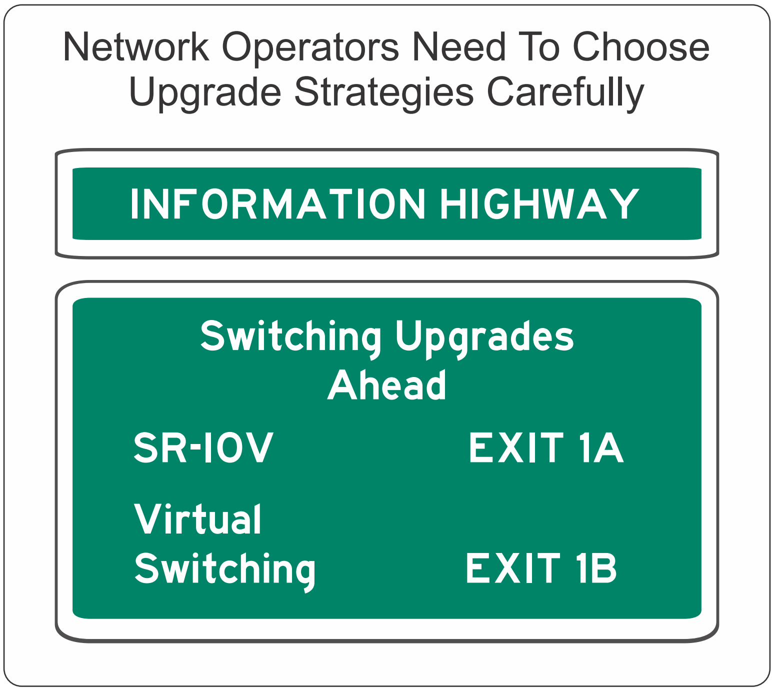
Network Operators Need To Choose Upgrade Strategies Carefully

This NFVI is in the process of being described within the NFV group at the European Technology Standards Institute, or ETSI. The purpose is to provide a framework of standards and technologies that network equipment vendors can rely on to ensure interoperability.

To maximize benefits of NFVI, some considerations must include:

- Predominance of VM-to-VM (East-West) communications over VM to physical network (North-South) communications
- Very high packet rates for some LTE nodes (in the range of millions of messages per second)
- Very high bandwidth for simplifying creation of managed virtual network operator (MVNO) infrastructures and other scenarios
- Tunneled telecom traffic such as GTP and mobile networks
- Overlay networks for service chaining and multi-tenancy

Virtual Networking is about virtualization of all components of what makes a networked infrastructure:

- Network cards
- Switches, Routers, Firewalls, Deep Packet Inspection (DPI)
- Wires

Operator costs are cut as previously expensive hardware solutions are being replaced with virtualized software switch solutions.

There are two broad approaches to networking virtualization in NFV Infrastructures:

- External specialized hardware: VMs communicate through virtual switch and router functions located in a dedicated hardware
- Embedded host software: VMs communicate through virtual switch and router functions located in the same host server

In other words, either the virtual networking is driven by evolved networking hardware or it is driven by software embedded in the commercial-off-the-shelf (COTS) platforms used for compute virtualization. When evaluating hardware versus software based switching solutions for virtual networking or NFVI, operators must evaluate options according to requirements that include:

- Portability: hardware independence
- Scalability: bandwidth and latency
- Performance: VM density and core count

This paper will show that all of these requirements and more can be met with accelerated software-

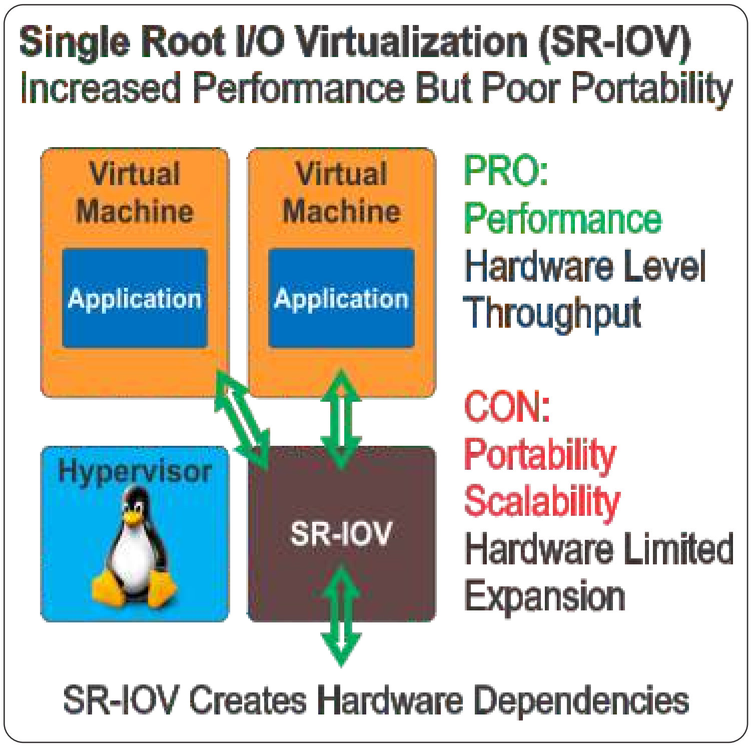
Single Root I/O Virtualization (SR-IOV)
Increased Performance But Poor Portability

SR-IOV Creates Hardware Dependencies

based virtual switching.

An Examination Of Hardware Switching

Until now, the service operator upgrade path was fairly well defined: buy proprietary hardware if, and when, the vendors make updates. The business case to be made needs to recoup capital expenditures with increased services in the shortest amount of time.

As a result, operators are often at the mercy of one or two large network vendors' equipment schedule, not to mention pricing.

In addition, it creates a much bigger issue of limiting the upgrade cycle based on availability versus demand. This can translate into inefficient services, outdated products and lost opportunities for the service operators.

Single Root I/O Virtualization (SR-IOV): Upgrade Solution, But At What Cost?

SR-IOV provides a mechanism by which a Single Root Function, such as a single Ethernet port, can appear to be multiple, separate devices.

Different applications and virtual machines can access the same port device as a unique affiliation. Many devices can share the same I/O port and access unique resources.

Other common SR-IOV benefits include high performance and reduced latency, due to direct access to hardware from virtual machine environments.

Capital and operational costs can be reduced:

- Power savings through the use of a single device for many virtual machines
- Reduced adapter count
- Less cabling
- Fewer switch ports

A unique problem this presents is a primary device dependence. VM to network throughput is improved but hardware independence is lost. Even if machines are virtualized to accommodate a single root source, there will be constraints on that device. The VMs run NIC-dependent drivers and the number of VMs in the system is limited by the number of virtual functions. Because of this hardware dependence with SR-IOV, VMs are not portable and upgrades can cause down-time.

Furthermore, VM-to-VM workloads bump into PCI Express limitations whereby only a fraction of the total bandwidth can flow from the NIC to the host. Therefore, bandwidth constraints may actually necessitate using a software switching based approach versus SR-IOV.

An Examination of Software Virtual Switching

Looking at modern virtualized network infrastructures, several immediate benefits are

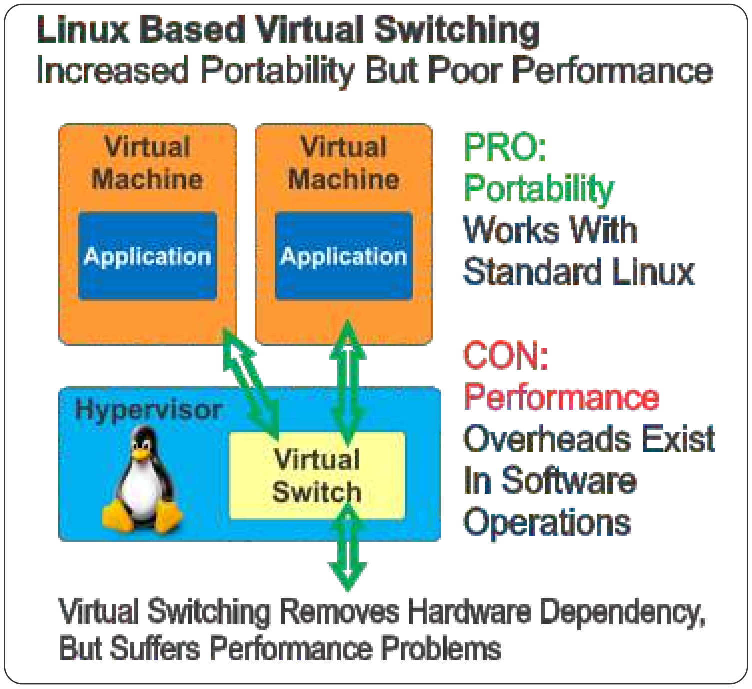
Linux Based Virtual Switching
Increased Portability But Poor Performance

Virtual Switching Removes Hardware Dependency, But Suffers Performance Problems obvious. We can see that previous hardware limitations can be resolved, if not eliminated, through strategic Linux-based software solutions.

With modern software architectures, the price structure has changed so much that hardware cost is no longer the issue. Bigger problems are physical space in the data center, power/cooling costs, and management.

Problems of software-based switching include the additional overhead to each I/O operation due to the abstraction layer between the guest driver and the I/O hardware. Until further work is done to abstract available acceleration mechanisms, software switching loses acceleration benefits. That said, proprietary virtual NICs can today offer some acceleration methods available in hardware.

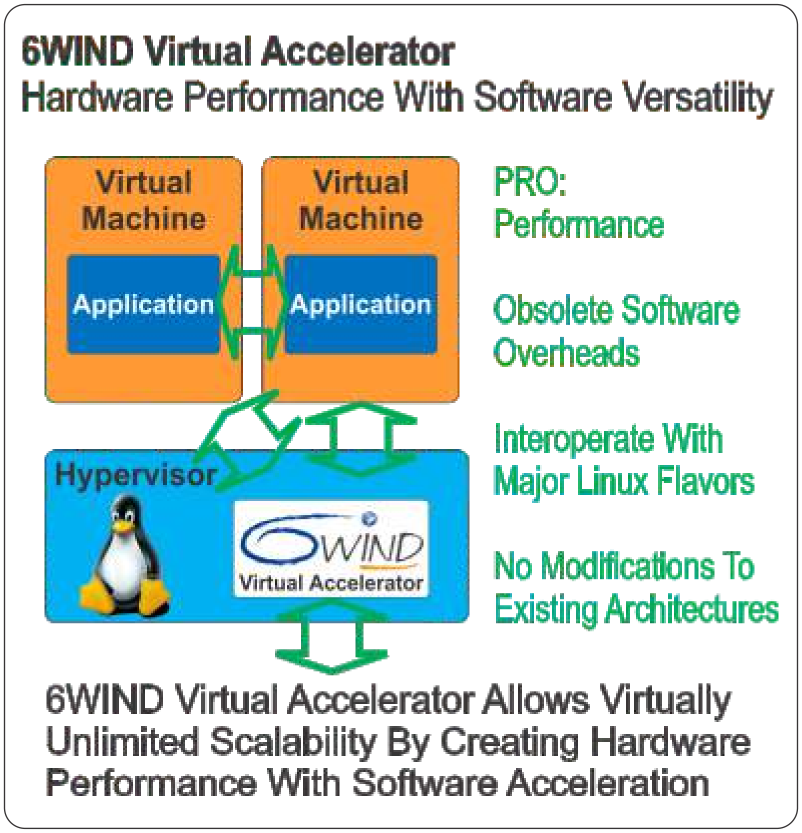
6WIND Virtual Accelerator
Hardware Performance With Software Versatility
Virtual Machine
Application
Virtual Machine
Application
Hypervisor
6WIND
Virtual Accelerator
PRO:
Performance
Obsolete Software Overheads
Interoperate With Major Linux Flavors
No Modifications To Existing Architectures
6WIND Virtual Accelerator Allows Virtually Unlimited Scalability By Creating Hardware Performance With Software Acceleration

Standard implementations based on a virtual switch, such as Open vSwitch or Linux bridging, do provide hardware independence but do not provide the right level of throughput for communication between VMs and between a VM and the outside world. This is due to bottlenecks in the Linux operating system, in the hypervisor, in the virtual switch implementation itself and in the communication channel between the hypervisor and the virtual machines.

This decrease in performance can be substantial, especially with cloud infrastructures that often need to scale dramatically.

Technology Summation

The transition from dated, centrally governed monolithic telecom infrastructures to a modern, distributed cloud-type network can be evaluated on the following criteria:

1. **Portability** is defined here as the ability to:

- install a virtual network function (VNF) on any hardware without modification
- migrate a VNF from one server to another one
- connect two or more VNFs located on any host without any VNF configuration

As described earlier, SR-IOV does not allow for portability due to hardware dependence. By virtualizing entire network functions, software switching allows complete hardware independence.

2. Scalability: VM Density and Core Count

Intuitively, VM density is higher with SR-IOV than with Linux-based software switching, because SR-IOV does not require switching cores.

However, in telecom networks, most of the traffic is tunneled using MPLS and GTP, and in data centers, overlay networks add another layer of encapsulation. As RSS (Receive Side Scaling) cannot load balance traffic between cores, it must be done by software. With software switching, it is the responsibility of one of the hypervisor's switching cores. With SR-IOV, this requires the dedication of additional VM cores to distribute traffic to other cores.

The bottom line is that the core count is identical with SR-IOV and software-based switching.

So there is no clear advantage of one solution over the other regarding density except when very high bandwidth and service chaining is required, such as in cloud infrastructures, in which case software switching presents an advantage for VM-to-VM communications.

3. Performance: Bandwidth and Latency

Network latency is directly proportional to both link bandwidth and number of hops within a data center. And overall latency is highly dependent on the application involved. However, when there is a high level of VM-to-VM traffic inside a single host, there is a clear advantage for the software approach as there are less hops than with a hardware approach.

Moreover, by providing well-defined APIs, software-based switching makes applications independent from the current limitation of acceleration hardware. When new features are supported by hardware (such as enhanced RSS, crypto, storage), they are transparently leveraged by applications, making the software-based approach future-proof.

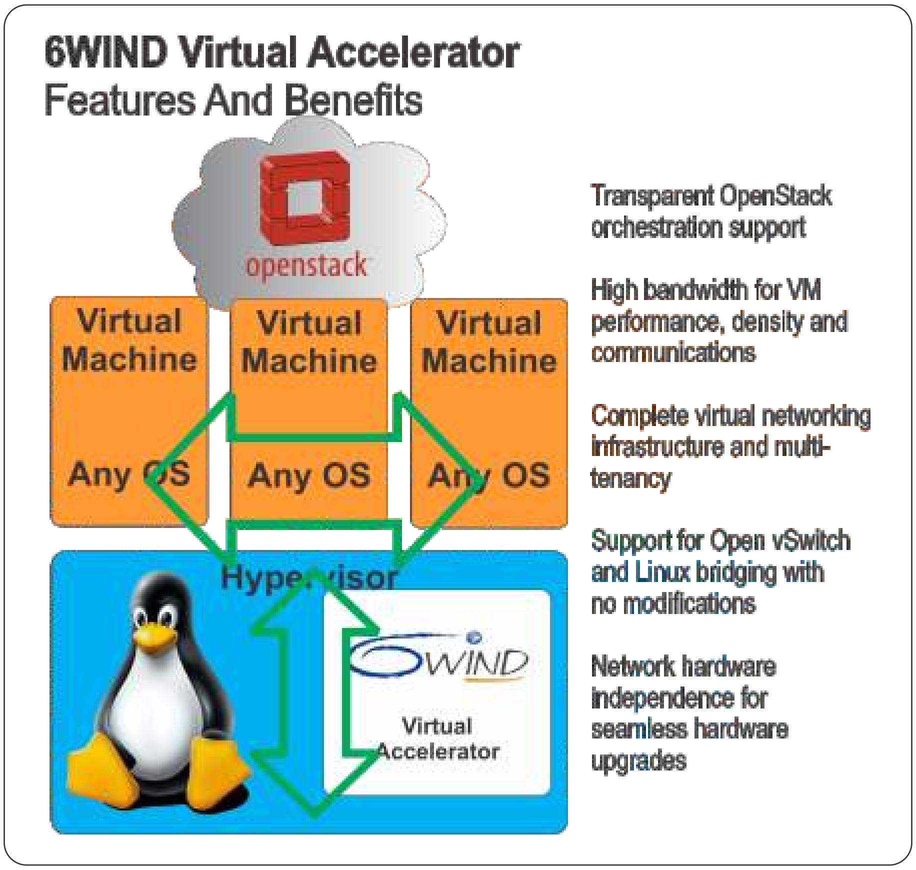
6WIND Virtual Accelerator
Features And Benefits
openstack
Virtual Machine
Any OS
Virtual Machine
Any OS
Virtual Machine
Any OS
Hypervisor
6WIND
Virtual Accelerator
Transparent OpenStack orchestration support
High bandwidth for VM performance, density and communications
Complete virtual networking infrastructure and multi-tenancy
Support for Open vSwitch and Linux bridging with no modifications
Network hardware independence for seamless hardware upgrades

Network Functions Virtualization enables traditionally hardware based functions to become virtualized and operate on commodity servers, and can overcome the limitations of SR-IOV thanks to software-based acceleration. Then to increase the scalability and performance of software-based switching, accelerated virtual switching solutions can be very effective. Otherwise, it will not be possible to maintain the performance of traditional hardware based solutions while gaining the desired promise of flexibility and cost savings.

6WIND Virtual Accelerator: Wire Speed Virtual Switching From Common Hardware

6WIND predicted the rapid availability of commodity hardware, and has created software applications specifically designed to bridge this performance gap of software versatility over hardware performance. 6WINDGate is a high performance networking stack that has been used widely for designing and building carrier grade infrastructure from commodity hardware, beginning with 3G/4G core infrastructure equipment and high performance network appliances. With 6WINDGate acceleration, performance gains of more than 10X are realized from standard Linux platforms.

Built on the success of 6WINDGate, 6WIND Virtual Accelerator provides accelerated virtual switching and networking features for virtual infrastructures to enable NFV. 6WIND Virtual Accelerator runs within the hypervisor domain with a hardware-independent architecture that allows new and existing VNFs to be integrated quickly onto COTS servers. As a transparent virtual infrastructure acceleration solution, Virtual Accelerator is provided as a simple software package that does not modify existing software such as Open vSwitch (OVS), Linux, Hypervisors and OpenStack.

6WIND Virtual Accelerator delivers:

- Network hardware independence for seamless hardware upgrades, including 10G to 40G to 100G ports
- Wire speed performance required to enable high density, compute intensive VMs on a single server
- Foundation for live migration of VMs over disparate hardware platforms
- Flexible virtual switching support for Open vSwitch and Linux bridging with no modifications
- Complete virtual networking infrastructure with VLAN, VXLAN, Virtual Routing, IP Forwarding, Filtering and NAT
- Native Virtio support for VMs based on different OSs
- High bandwidth for VM-to-VM communications required for Service Chaining
- Transparent orchestration support for OpenStack