



Global Knowledge®



Expert Reference Series of White Papers

What's New in VMware vSphere™ 4: Virtual Networking

What's New in VMware vSphere™ 4: Virtual Networking

Introduction

VMware vSphere™ introduces a number of new features and capabilities to virtual networking under VMware® vNetwork. vNetwork is the new name to describe the collection of networking technologies for optimally integrating networking and I/O functionality into vSphere.

These vNetwork enhancements provide the server admin and network admin with an unprecedented level of control while simplifying deployment, ongoing management, and troubleshooting.

This paper provides an overview of the major enhancements introduced with VMware vNetwork.

VMware vNetwork: Summary of Enhancements

The major enhancements to VMware vNetwork are as follows. These are further explained in the sections below.

- **vNetwork Distributed Switch (vDS)**—VMware's next generation virtual networking solution for spanning multiple hosts with a single virtual switch representation. vDS enables and includes some additional enhancements as follows:
 - Private VLANs
 - Network VMotion—tracking of VM networking state, improving troubleshooting and enabling
 - 3rd Party Virtual Switch support with the Cisco Nexus 1000V Series Virtual Switch
 - Bi-directional traffic shaping
- **VMXNET3**—Third generation para-virtualized NIC
- **IPv6**—support extended to vmkernel and Service Console ports

vNetwork Distributed Switch

The vNetwork Distributed Switch (vDS) extends the features and capabilities and features of virtual networks while simplifying provisioning and the ongoing process of configuration, monitoring, and management.

With ESX 3.5 and prior releases, virtual networks were constructed using virtual switches or vSwitches. Each ESX host would use one or more vSwitches to connect the VMs with the server NICs and the outside physical network.

Simplified Network Provisioning, Configuration and Management with vDS

In addition to continuing support for the vSwitch (now known as the Standard Switch), vSphere introduces an additional choice for VMware virtual networking with the vNetwork Distributed Switch. vDS eases the management burden of per host, virtual switch configuration management by treating the network as an aggregated resource. Individual, host-level virtual switches are abstracted into a single large vNetwork Distributed Switch that spans multiple hosts at the Datacenter level. Port Groups become Distributed Virtual Port Groups (DV Port Groups) that span multiple hosts and ensure configuration consistency for VMs and virtual ports necessary for such functions as VMotion.

Figures 1 and 2 illustrate the conceptual difference in management for a Standard Switch environment versus a vDS environment. Each of the Standard Switches in Figure 1 requires a separate configuration from a separate management panel. The vDS in Figure 2 requires just one management panel for the single switch that spans multiple hosts.

Figure 1 - Standard Switches are individually managed and configured.





Figure 2 - Management of vNetwork Distributed Switches is independent of the number of hosts.



Distributed Virtual Port Groups and Distributed Virtual Uplinks

Many of the concepts involved in configuring and managing a Standard Switch are carried forward with the vDS.

Distributed Virtual Port Groups (DV Port Groups) are port groups associated with a vDS and specify port configuration options for each member port. DV Port Groups define how a connection is made through the vDS to the Network. Configuration parameters are similar to those available with Port Groups on Standard Switches. The VLAN ID, traffic shaping parameters, port security, teaming and load balancing configuration, and other settings are configured here.

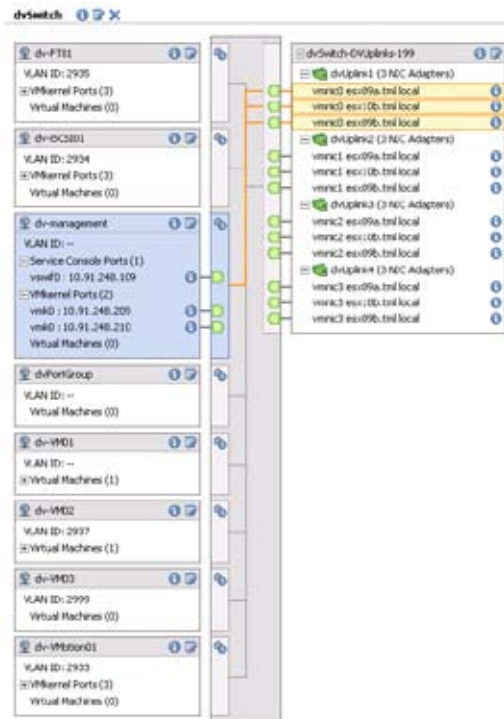
Distributed Virtual Uplinks (dvUplinks) are a new concept introduced with vDS. dvUplinks provide a level of abstraction for the physical NICs (vmnics) on each host. NIC teaming, load balancing, and failover policies on the vDS and DV Port Groups are applied to the dvUplinks and not the vmnics on individual hosts. Each vmnic on each host is mapped to a dvUplink, permitting teaming and failover consistency irrespective of vmnic assignments. This is illustrated in the dvUplink box in Figure 3. vmnic0 on each of the three hosts (esx09a, esx10b, esx9b) is mapped to dvUplink1. If desired, any of the vmnics could be assigned on any of the hosts to dvUplink1.

Figure 3 illustrates the vDS view from a vSphere client for a three host sample environment.

New Features with vDS

In addition to easing the configuration and management burden, vDS brings with it a number of new features and capabilities to address some common and emerging virtual network requirements. Note that these features are not available with Standard Switches.

Figure 3 – An example vDS for a small three host environment showing highlighted path through switch to dvUplinks for the dv-management Distributed Virtual Port Group.



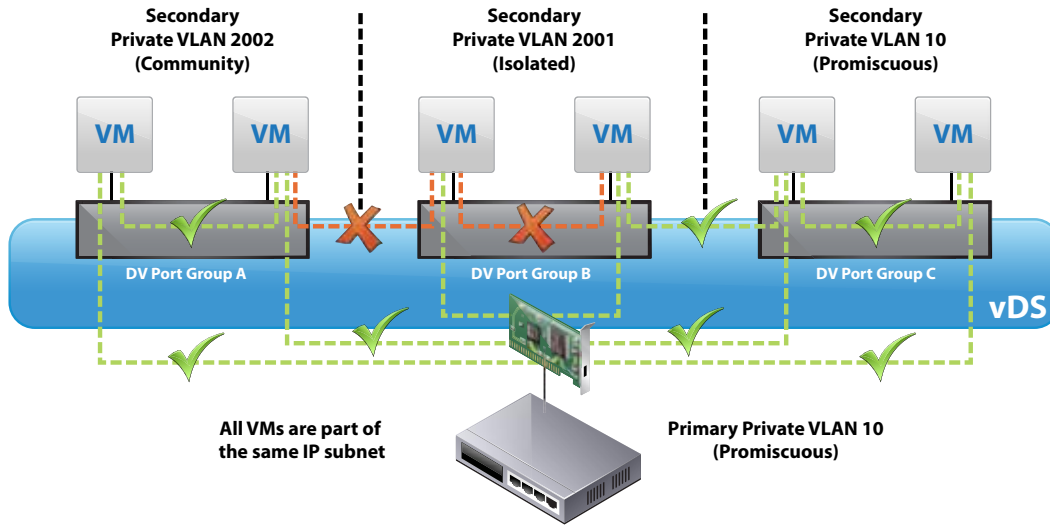
Private VLANs

Private VLAN (PVLAN) support enables broader compatibility with existing networking environments using Private VLAN technology. Private VLANs enable users to restrict communication between virtual machines on the same VLAN or network segment, significantly reducing the number of subnets needed for certain network configurations.

Figure 4 illustrates how this concept works with a vDS. Private VLANs are configured on a vDS with allocations made to the Promiscuous Private VLAN, the Community Private VLAN and the Isolated Private VLAN. DV Port Groups can then use one of these Private VLANs and VMs are then assigned to a DV Port Group. Within the subnet, VMs on the Promiscuous Private VLAN can communicate with all VMs; VMs on the Community Private VLAN can communicate amongst themselves and those on the Promiscuous Private VLAN; VMs on the isolated Private VLAN can only communicate with VMs on the Promiscuous Private VLAN.

Note that the adjacent physical switches must support Private VLANs and be configured to support the Private VLANs allocated on the vDS.

Figure 4 - Private VLANs provide a simple way of selectively isolating VMs without exhausting IP subnets.



Network VMotion

Network VMotion is the tracking of virtual machine networking state (e.g. counters, port statistics) as the VM moves from host to host on a vNetwork Distributed Switch. This provides a consistent view of a virtual network interface regardless of the VM location or VMotion migration history. This greatly simplifies network monitoring and troubleshooting activities where VMotion is used to migrate VMs between hosts.

Bi-directional Traffic Shaping

vDS expands upon the egress only traffic shaping feature of Standard Switches with bi-directional traffic shaping capabilities. Egress (from VM to network) and now ingress (from network into VM) traffic shaping policies can now be applied on DV Port Group Definitions.

Traffic shaping is useful in cases where you may wish to limit the traffic to or from a VM or group of VMs to either protect a VM or other traffic in an oversubscribed network.

Policies are defined by three characteristics: average bandwidth, peak bandwidth, and burst size. See Figure 5 below.

Figure 5 - Traffic shaping policy definition on DV Port Group.

Policies	
Ingress Traffic Shaping	
Status:	Enabled
Average Bandwidth:	80000 Kbits/sec
Peak Bandwidth:	100000 Kbits/sec
Burst Size:	102400 Kbytes
Egress Traffic Shaping	
Status:	Enabled
Average Bandwidth:	50000 Kbits/sec
Peak Bandwidth:	100000 Kbits/sec
Burst Size:	102400 Kbytes

Third Party Virtual Switch Support with the Cisco Nexus 1000V Series Virtual Switch

The vNetwork Distributed Switch includes switch extensibility for seamless integration of 3rd party control planes, data planes, and user interfaces. Cisco has collaborated with VMware to exploit this extensibility to produce the Cisco Nexus 1000V Series Virtual Switch.

The Cisco Nexus 1000V uses the same distributed switching model as the VMware vNetwork Distributed Switch. Virtual Ethernet Modules (VEMs) are the switching data planes on each ESX host and provide the frame forwarding capabilities. The VEMs leverage the ESX host APIs and so can leverage the same physical NICs and HCL (Hardware Compatibility List) as the VMware Standard Switch and vNetwork Distributed Switch. Virtual Supervisor Modules (VSMs) are implemented on the Cisco NX-OS operating system. They provide the control plane function for the VEMs and can exist as a guest VM or standalone appliance.

VSMs provide a familiar Cisco CLI (Command Line Interface) for management and configuration. They also communicate with vCenter Server for optional management and configuration through a vSphere Client.

The Cisco Nexus 1000V has an expanded feature set similar to that provided by physical Cisco Catalyst and Nexus switches.

For more information on the Cisco Nexus 1000V, go to <http://cisco.com/go/nexus1000v>.

Additional Features Introduced with VMware vNetwork

VMXNET3

VMXNET3 builds upon VMXNET and Enhanced VMXNET as the third generation paravirtualized virtual networking NIC for guest operating systems.

New VMXNET3 features over previous version of Enhanced VMXNET include:

- MSI/MSI-X support (subject to guest operating system kernel support)
- Receive Side Scaling (supported in Windows 2008 when explicitly enabled through the device's Advanced configuration tab)
- IPv6 checksum and TCP Segmentation Offloading (TSO) over IPv6
- VLAN off-loading
- Large TX/RX ring sizes (configured from within the virtual machine)

IPv6

IPv6 (IP version 6) is the successor to the dominant IPv4 protocol used in the Internet today. IPv6 incorporates a number of improvements over IPv4, namely integrated network security, plus an increased address space to alleviate IPv4 address exhaustion.

IPv6 support for guest operating systems was introduced in VMware ESX 3.5. With vSphere, IPv6 support is extended to include the vmkernel and service console allowing IP storage and other ESX services to communicate over IPv6.

VMDirectPath

VMDirectPath is a new capability provided in vSphere for direct assignment of PCI devices to a VM for guest control of physical hardware.

VMDirectPath is designed for special purpose I/O appliances and high performance VMs that require the portability and management benefits of a VM, but do not need support for additional VM functions such as VMotion, fault tolerance and suspend/resume.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge. Check out the following Global Knowledge courses:

[VMware vSphere: Install, Configure, Manage \[V4\]](#)

[VMware vSphere: Fast Track \[V4\]](#)

[ICNX1 - Implementing and Configuring the Nexus 1000V](#)

For more information or to register, visit www.globalknowledge.com or call **1-800-COURSES** to speak with a sales representative.

Our courses and enhanced, hands-on labs offer practical skills and tips that you can immediately put to use. Our expert instructors draw upon their experiences to help you understand key concepts and how to apply them to your specific work situation. Choose from our more than 700 courses, delivered through Classrooms, e-Learning, and On-site sessions, to meet your IT and management training needs.

VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com
Copyright © 2009 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMW_09Q1_WP_vSphereNetworking_P8_R1