



What's New in VMware® vSphere™ 4.1 — Networking

VMware vSphere 4.1

WHITE PAPER

Introduction

VMware® vSphere™ 4.1 (“vSphere”) introduces a number of enhancements and new features to virtual networking.

- **Network I/O Control (NetIOC)**—flexibly partition and assure service for ESX/ESXi traffic types and flows on a vNetwork Distributed Switch (vDS)
- **Load-Based Teaming (LBT)**—an additional and selectable load-balancing policy on the vDS to enable dynamic adjustment of the load distribution over a team of NICs
- **Network performance**—vmkernel TCP/IP stack and guest virtual-machine network performance enhancements
- **Scale**—enhancements to network scaling with the vDS
- **IPv6 NIST Compliance**—IPv6 enhancements to comply with U.S. National Institute of Standards and Technology (NIST) Host Profile
- **Cisco Nexus 1000V Enhancements**—support for new features and enhancements on the Cisco Nexus 1000V

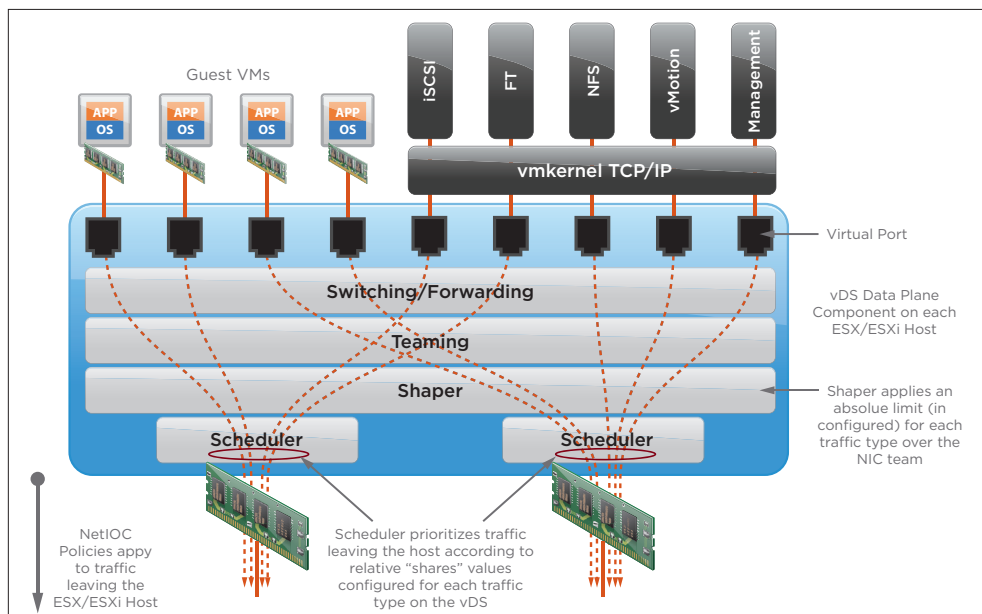
Network I/O Control

Network I/O Control (NetIOC) is a new traffic-management feature of the vDS. NetIOC implements a software scheduler within the vDS to isolate and prioritize specific traffic types contending for bandwidth on the uplinks connecting ESX/ESXi 4.1 hosts with the physical network.

NetIOC is able to individually identify and prioritize the following traffic types leaving an ESX/ESXi host on a vDS-connected uplink:

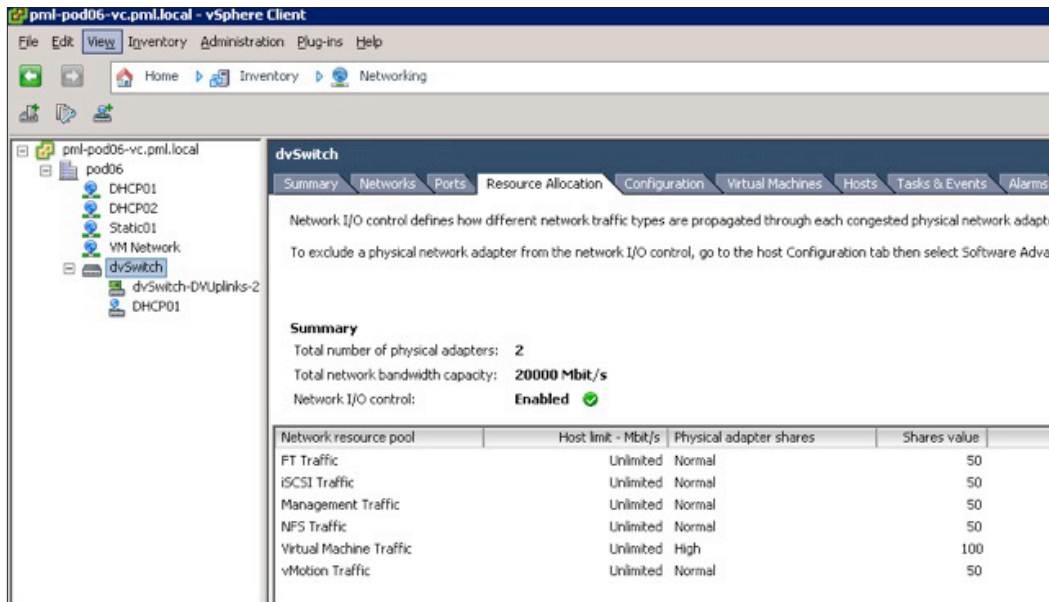
- virtual-machine traffic
- management traffic
- iSCSI
- NFS
- VMware Fault Tolerance (VMware FT) logging
- VMotion™

NetIOC is particularly applicable to environments in which multiple traffic types are converged over a pair of 10GbE interfaces. If an interface is oversubscribed (that is, more than 10Gbps is contending for a 10GbE interface), NetIOC is able to ensure each traffic type is given a selectable and configurable minimum level of service.



Configuring NetIOC with shares and limits

NetIOC uses two parameters to prioritize and limit the network traffic leaving a vDS—shares and limits. These values are configured through the vSphere client on the vCenter Server (see figure).



Shares—The shares value specifies the relative importance of a traffic type scheduled for transmission on a physical NIC (vmnic). Shares are specified in abstract units between 1 and 100. Bandwidth for the link is divided among the traffic types according to their relative shares value.

For example, assume we have a 10GbE link, with virtual-machine traffic set to 100 shares, vMotion traffic set to 50 shares, and VMware FT logging traffic set to 50 shares. If virtual-machine traffic and vMotion traffic are both contending for the 10GbE link, then the virtual-machine traffic (100 shares) will get 67% (670Mbps) of the link, and vMotion (50 shares) will get 33% (330Mbps) of the link. If we have all three of these traffic types active and contending for the link, then virtual-machine traffic (100 shares) will get 50% (500Mbps), vMotion (50 shares) will get 25% (250Mbps), and VMware FT logging (50 shares) will get 25% (250Mbps). If no other traffic types are contending for the link at that moment, then each traffic type can consume the entire link (or up to the host limit, if set).

Limits—The limits value specifies an absolute maximum limit on egress traffic for that traffic type on a host. Limits are specified in Mbps. The limit is an aggregate for that traffic type and applies regardless of the number of physical NICs (vmnics) in the NIC team.

NOTE: Limits are applied to the network traffic before the shares. Limits apply over a team, while shares schedule and prioritize traffic for each physical NIC.

Load-Based Teaming (LBT)

LBT is another traffic-management feature of the vDS introduced with vSphere 4.1. LBT avoids network congestion on the ESX/ESXi host uplinks caused by imbalances in the mapping of traffic to those uplinks.

LBT enables customers to optimally use and balance network load over the available physical uplinks attached to each ESX/ESXi host. LBT helps avoid situations where one link may be congested, while other links may be relatively underused.

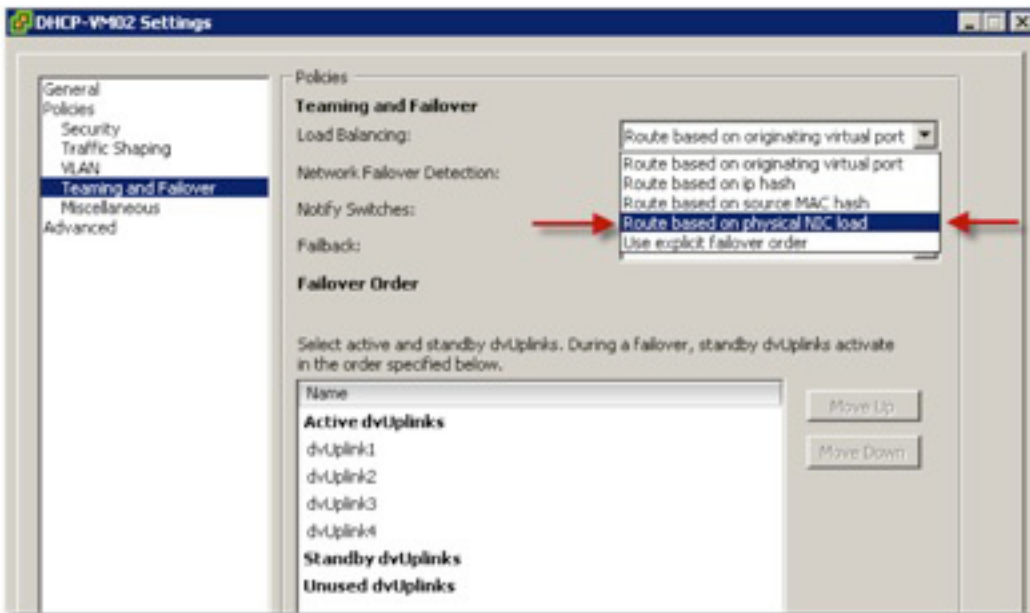
How LBT works

LBT dynamically adjusts the mapping of virtual ports to physical NICs to best balance the network load entering or leaving the ESX/ESXi 4.1 host. When LBT detects an ingress- or egress- congestion condition on an uplink, signified by a mean utilization of 75% or more over a 30-second period, it will attempt to move one or more of the virtual ports to vmnic-mapped flows to lesser-used links within the team.

Configuring LBT

LBT is an additional load-balancing policy available within the teaming and failover of a dvPortGroup on a vDS. LBT appears as the "Route based on physical NIC load."

LBT is not available on the vNetwork Standard Switch (vSS).



Network performance

vSphere 4.1 incorporates a number of network performance enhancements that affect native guest virtual-machine throughput as well as that for vmkernel-based ESX/ESXi applications such as vMotion, VMware FT Logging, NFS and iSCSI. Some of these improvements are as follows (note that observed performance increases will vary according to the platform and other outside factors):

- vMotion throughput—this is increased, leading to a 50% time reduction to migrate a virtual machine in many cases.
- vMotion concurrency—vSphere will automatically increase the maximum number of concurrently allowed vMotion to eight (up from a maximum of two with vSphere 4.0) when 10GbE links are employed.
- NFS—increased throughput for both read and write.
- Native virtual machine throughput increases.

Scaling Increases

Many of the networking maximums are raised in vSphere 4.1. Some of the notable increases are as follows:

- 350 hosts per vDS (64 in vSphere 4.0)
- 20,000 ports per vDS (4,096 in vSphere 4.0)
- 32 vDS per vCenter (16 in vSphere 4.0)

The networking maximums are subject to change. For the complete list of maximums, consult the “Networking Maximums” section in the *Configuration Maximums* document for vSphere 4.1 available at www.vmware.com/support/pubs.

IPv6 U.S. NIST Compliance

vSphere 4.1 is undergoing testing for U.S. NIST Host Profile compliance that includes requirements for IPsec and IKEv2 functionality (with the exception of MLDv2 plus PKI and DH-24 support within IKE).

NOTE: *The NIST Profiles are described in <http://www.antd.nist.gov/usgv6/usgv6-v1.pdf>.*

In vSphere 4.1, IPv6 is supported for:

- Guest virtual machines
- ESX/ESXi management
- vSphere client
- vCenter Server
- vMotion
- IP storage (iSCSI, NFS)—experimental

NOTE: *IPv6 is not supported for vSphere vCLI, VMware HA and VMware FT logging. IKEv2 is disabled by default.*

Cisco Nexus 1000V

A new release of the Cisco Nexus 1000V with new and improved features and functionality is available with vSphere 4.1. For more details, refer to the Nexus 1000V site at www.cisco.com/go/nexus1000v.

For More Information

VMware networking technology

<http://www.vmware.com/go/networking>

Networking blog

<http://www.blogs.vmware.com/networking>

VMware/Cisco partner site

<http://www.vmware.com/cisco>

