What's New in VMware vSphere™ 5.0

Networking

TECHNICAL MARKETING DOCUMENTATION
V 1.0/UPDATED APRIL 2011

vmware®

Table of Contents

Introduction

With the release of VMware vSphere™ 5.0 ("vSphere"), VMware brings a number of powerful new features and enhancements to the networking capabilities of the vSphere platform. These new network capabilities enable customers to run business-critical applications with confidence and provide the flexibility to enable customers to respond to business needs more rapidly. All the networking capabilities discussed in this document are available only with the VMware vSphere Distributed Switch (Distributed Switch).

There are two broad types of networking capabilities that are new or enhanced in the VMware vSphere 5.0 release. The first type improves the network administrator's ability to monitor and troubleshoot virtual infrastructure traffic by introducing features such as

- NetFlow
- Port mirror

The second type focuses on enhancements to the network I/O control (NIOC) capability first released in vSphere 4.1. These NIOC enhancements target the management of I/O resources in consolidated I/O environments with 10GB network interface cards. The enhancements to NIOC enable customers to provide end-to-end quality of service (QoS) through allocating I/O shares for user-defined traffic types as well as tagging packets for prioritization by external network infrastructure. The following are the key NIOC enhancements:

- User-defined resource pool
- vSphere replication traffic type
- IEEE 802.1p tagging

The following sections will provide higher-level details on new and enhanced networking capabilities in vSphere 5.0.

Network Monitoring and Troubleshooting

In a vSphere 5.0 environment, virtual network switches provide connectivity for virtual machines running on VMware® ESXi™ hosts to communicate with each other as well as connectivity to the external physical infrastructure. Network administrators want more visibility into this traffic that is flowing in the virtual infrastructure. This visibility will help them monitor and troubleshoot network issues. VMware vSphere 5.0 introduces two new features in the Distributed Switch that provide the required monitoring and troubleshooting capability to the virtual infrastructure.

NetFlow

NetFlow is a networking protocol that collects IP traffic information as records and sends them to a collector such as CA NetQoS for traffic flow analysis. VMware vSphere 5.0 supports NetFlow v5, which is the most common version supported by network devices. NetFlow capability in the vSphere 5.0 platform provides visibility into virtual infrastructure traffic that includes

- Intrahost virtual machine traffic (virtual machine–to–virtual machine traffic on the same host)
- Interhost virtual machine traffic (virtual machine–to–virtual machine traffic on different hosts)
- Virtual machine–physical infrastructure traffic

Figure 1 shows a Distributed Switch configured to send NetFlow records to a collector that is connected to an external network switch. The blue dotted line with arrow indicates the NetFlow session that is established to send flow records for the collector to analyze.

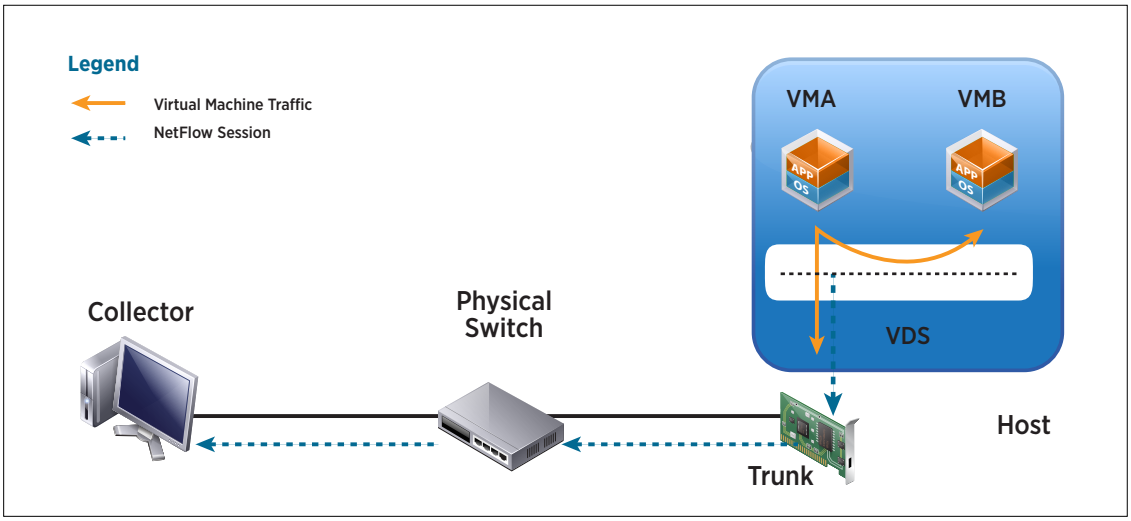

Figure 1. NetFlow Traffic

Usage

NetFlow capability on a Distributed Switch along with a NetFlow collector tool helps monitor application flows and measures flow performance over time. It also helps in capacity planning and ensuring that I/O resources are utilized properly by different applications, based on their needs.

IT administrators who want to monitor the performance of application flows running in the virtualized environment can enable flow monitoring on a Distributed Switch.

Configuration

NetFlow on Distributed Switches can be enabled at the port group level, at an individual port level or at the uplink level. When configuring NetFlow at the port level, administrators should select the NetFlow override tab, which will make sure that flows are monitored even if the port group–level NetFlow is disabled.

Port Mirror

Port mirroring is the capability on a network switch to send a copy of network packets seen on a switch port to a network monitoring device connected to another switch port. Port mirroring is also referred to as Switch Port Analyzer (SPAN) on Cisco switches. In VMware vSphere 5.0, a Distributed Switch provides a similar port mirroring capability to that available on a physical network switch. After a port mirror session is configured with a destination—a virtual machine, a vmknic or an uplink port—the Distributed Switch copies packets to the destination. Port mirroring provides visibility into

- Intrahost virtual machine traffic (virtual machine–to–virtual machine traffic on the same host)
- Interhost virtual machine traffic (virtual machine–to–virtual machine traffic on different hosts)

Figure 2 shows different types of traffic flows that can be monitored when a virtual machine on a host acts as a destination or monitoring device. All traffic shown by the orange dotted line with arrow is mirrored traffic that is sent to the destination virtual machine.

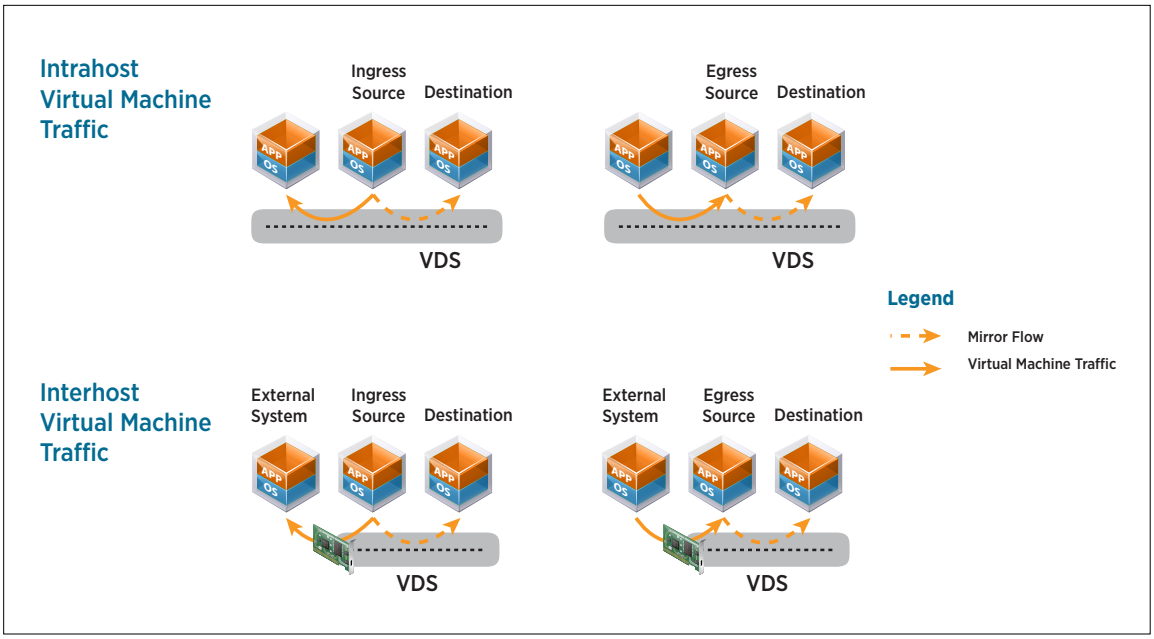

Figure 2. Port Mirror Traffic Flows When Destination Where Packets Are Mirrored Is a Virtual Machine

Usage

The port mirroring capability on a Distributed Switch is a valuable tool that helps network administrators in debugging network issues in a virtual infrastructure. The granular control over monitoring ingress, egress or all traffic of a port helps administrators fine-tune what traffic is sent for analysis.

Configuration

Port mirror configuration can be done at the Distributed Switch level, where a network administrator can create a port mirror session by identifying the traffic source that needs monitoring and the traffic destination where the traffic will be mirrored. The traffic source can be any port with ingress, egress or all traffic selected. The traffic destination can be any virtual machine, vmknic or uplink port.

Network Management and Configuration

The number of network devices, virtual and physical, in a datacenter is increasing. For a network or IT administrator, it becomes difficult to manually configure and monitor each of these devices. In such complex environments, with a large number of devices, configuration-related mistakes are very difficult to rectify. To solve this configuration problem, Cisco introduced Cisco Discovery Protocol (CDP). The VMware vSphere platform supported this protocol and helped simplify and automate the configuration process with Cisco network infrastructure.

With the release of vSphere 5.0, VMware now supports IEEE 802.1AB standard–based Link Layer Discovery Protocol (LLDP). LLDP helps management and configuration of heterogeneous network devices from different vendors.

LLDP

Discovery protocol is a data link layer protocol used to discover capabilities of network devices. LLDP is a standard-based (IEEE 802.1AB), vendor-neutral discovery protocol.

Usage

LLDP enables administrators to automate the deployment and configuration process in a complex network-switching environment. It also helps avoid downtime due to misconfiguration of network devices.

Configuration

This feature can be enabled at the Distributed Switch level by selecting either the CDP or LLDP discovery protocol type. Customers also can choose the mode of operation for the discovery protocol from the following three options:

- Listen
- Advertise
- Both

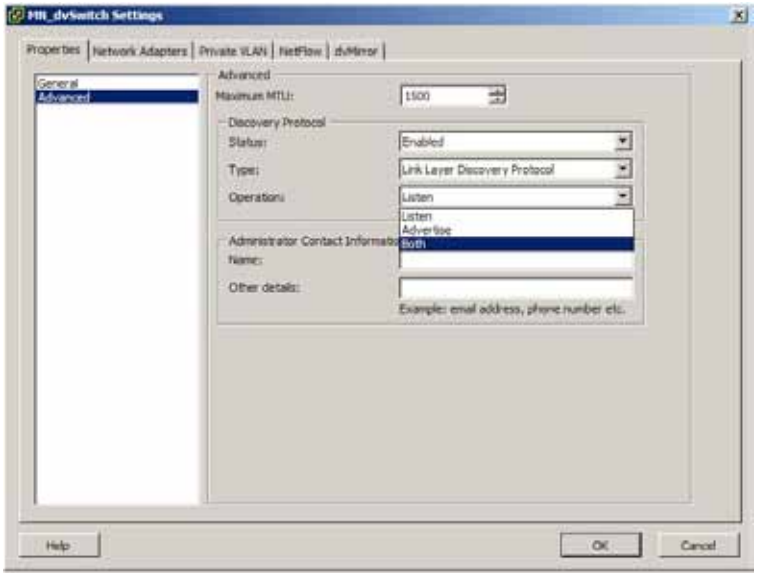

Figure 3. LLDP Configuration

Network I/O Control Enhancements

Consolidated I/O or I/O virtualization delivers similar benefits as provided by x86 virtualization in terms of better utilization and consolidation of resources. However, as multiple traffic types flow through a single physical network interface, it becomes important to manage the traffic effectively such that critical application flows don't suffer because of a burst of low-priority traffic. Network traffic management provides the required control and guarantee to different traffic types in the consolidated I/O environment. In the VMware vSphere 5.0 platform, NIOC supports traffic management capabilities for the following traffic types:

- Virtual machine traffic
- Management traffic
- iSCSI traffic
- NFS traffic
- Fault-tolerant traffic
- VMware vMotion™ traffic
- User-defined traffic
- vSphere replication traffic

Similar to CPU and memory resource allocation in the vSphere platform, a network administrator through NIOC can allocate I/O shares and limits to different traffic types, based on their requirements. In this new release of vSphere, NIOC capabilities are enhanced such that administrators can now create user-defined traffic types and

allocate shares and limits to them. Also, administrators can provide I/O resources to the vSphere replication process by assigning shares to vSphere replication traffic types. The following section describes in detail the two new traffic types introduced in vSphere 5.0.

User-Defined Network Resource Pools

User-defined network resource pools in vSphere 5.0 provide an ability to add new traffic types beyond the standard system traffic types that are used for I/O scheduling.

Figure 4 shows an example of a user-defined resource pool with shares, limits and IEEE 802.1p tag parameters described in a table. In this example, Tenant 1 and Tenant 2 are two user-defined resource pools with virtual machines connected to their respective independent port groups. Tenant 1, with three virtual machines, has five I/O shares. Tenant 2, with one virtual machine, has 15 I/O shares. This indicates that during contention scenarios, Tenant 2 virtual machines will have a higher guaranteed share than Tenant 1 virtual machines.

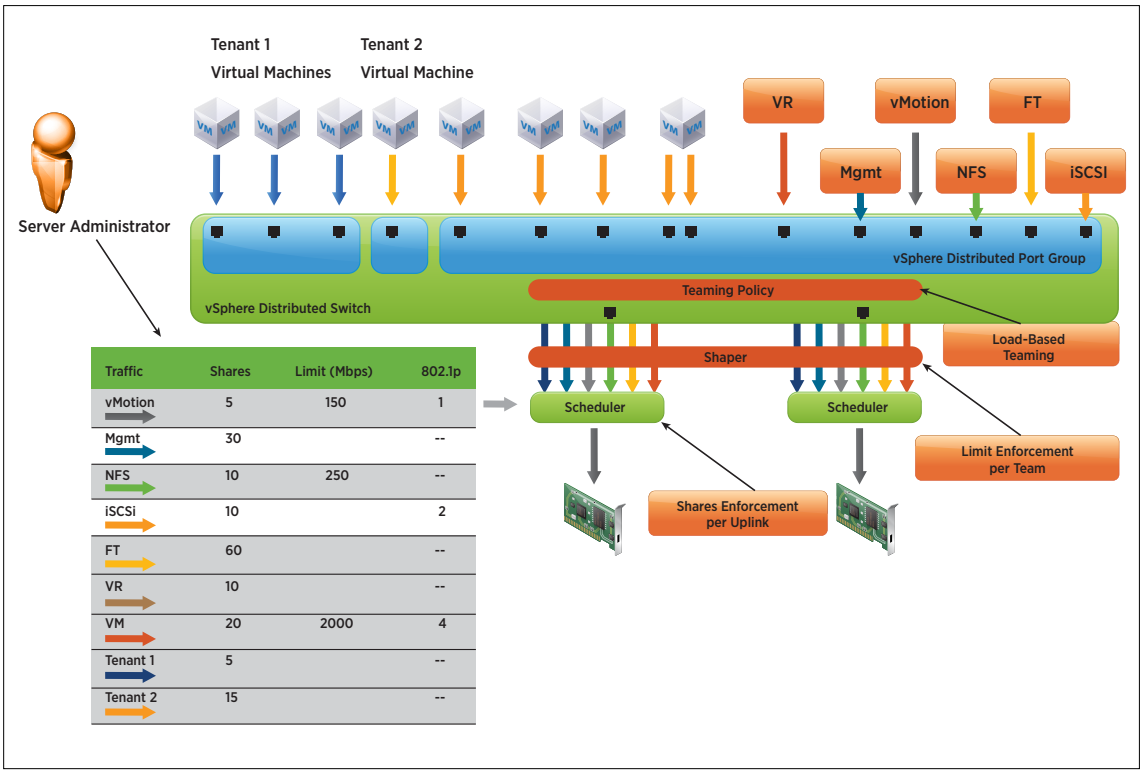

Figure 4. NIOC User-Defined Resource Pools: Tenant 1 and Tenant 2

Usage

When customers are deploying critical applications on virtual infrastructure, they can utilize this advanced feature to reserve I/O resources for the important, business-critical application traffic and provide SLA guarantees.

Service providers who are deploying public clouds and serving multiple tenants can now define and provision I/O resources per tenant, based on each tenant's need.

Configuration

The new resource pools can be defined at the Distributed Switch level by selecting the resource allocation tab and clicking on new network resource pools. After a new network resource pool is defined with shares and limits parameters, that resource pool can be associated with a port group. This association of a network resource pool with a port group enables customers to allocate I/O resources to a group of virtual machines or workloads.

vSphere Replication Traffic

vSphere replication is a new system traffic type that carries replication traffic from one host to another. NIOC now supports this new traffic type along with other system and user-defined traffic types.

Usage

Customers implementing a disaster recovery (DR) solution with VMware vCenter Site Recovery Manager (Site Recovery Manager) and vSphere replication can use this vSphere replication traffic type to provide required network resources to the replication process.

Configuration

A vSphere replication traffic type can be configured on a Distributed Switch under the resource allocation tab. This traffic type is now part of the system network resource pool. Customers can allocate shares and limits parameters to this traffic type.

IEEE 802.1p Tagging

IEEE 802.1p is a standard for enabling QoS at MAC level. The IEEE 802.1p tag provides a 3-bit field for prioritization, which allows packets to be grouped into seven different traffic classes. The IEEE doesn't mandate or standardize the use of recommended traffic classes. However, higher-number tags typically indicate critical traffic that has higher priority. The traffic is simply classified at the source and sent to the destination. The layer-2 switch infrastructure between the source and destination handles the traffic classes according to the assigned priority. In the vSphere 5.0 release, network administrators now can tag the packets going out of the host.

Usage

Customers who are deploying business-critical applications in a virtualized environment now have the capability to guarantee I/O resources to these workloads on the host. However, it is not sufficient to provide I/O resources just on the host. Customers must think about how to provide end-to-end QoS to the business-critical application traffic. The capability of a Distributed Switch to provide an IEEE 802.1p tag helps such customers meet those requirements for end-to-end QoS or service-level agreements.

Configuration

IEEE 802.1p tagging can be enabled per traffic type. Customers can select the Distributed Switch and then the resource allocation tab to see the different traffic types, including system and user-defined traffic types. After selecting a traffic type, the user can edit the QoS priority tag field by choosing any number from 1 to 7. Figure 5 is the screenshot of QoS priority tag configuration for the *MyVMTraffic* traffic type.

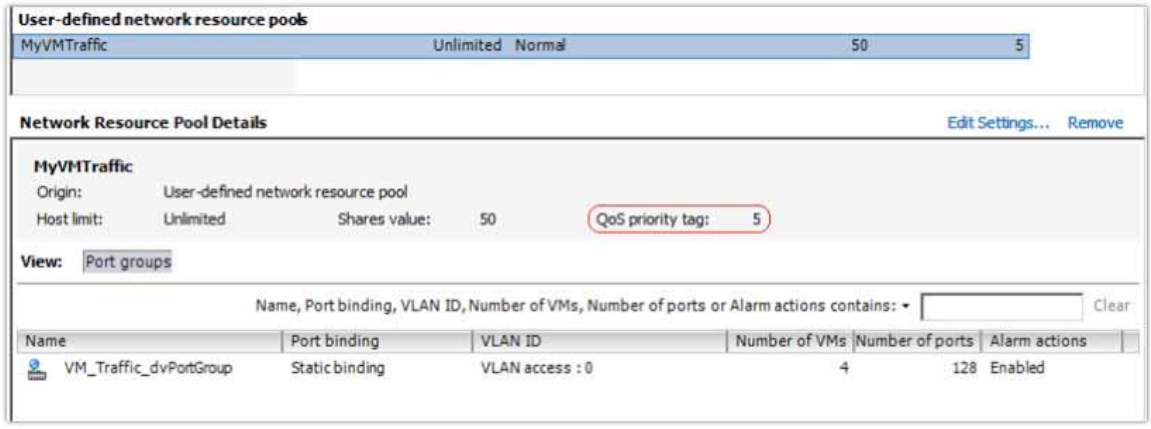

Figure 5. Configured QoS Priority for User-Defined *MyVMTraffic* Resource Pool

Conclusion

The VMware vSphere 5.0 release builds on the advanced networking features supported in vSphere 4.1. The new and enhanced features in vSphere 5.0 provide additional tools and capabilities that will simplify the lives of network administrators and provide more flexibility in deploying a cloud infrastructure. Features such as NetFlow and port mirror provide network administrators with a visibility into virtual infrastructure traffic through the familiar tools and protocols. Enhanced network I/O control helps customers manage I/O resource pools, based on workload and system resource demand, and also provides end-to-end service-level agreements for critical applications.

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com