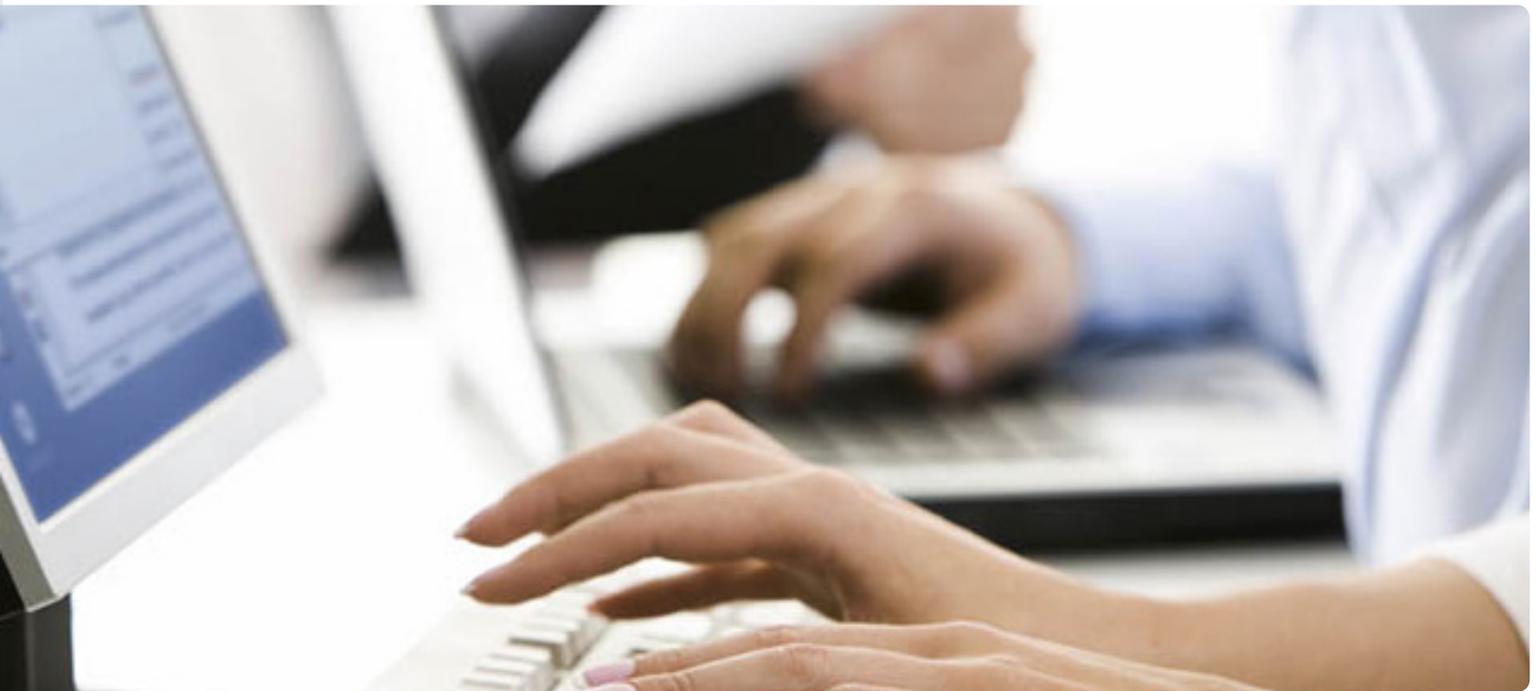


# Top seven networking concepts every virtualization administrator should understand

Written by Scott D. Lowe, Founder and Managing Consultant, The 1610 Group



## Abstract

Networking in a virtual environment is different from networking in the physical world—but no less critical. Because the network remains the foundation of the IT infrastructure, virtualization administrators need to understand its nature and behavior. This paper explores seven critical networking concepts that virtualization administrators must understand in order to maximize system performance while controlling costs.

## Introduction

Virtualization administrators need to know a lot, but networking remains a key element.

A long time ago in an IT department far away, there used to exist three separate teams in the IT department that handled infrastructure activities. These teams—systems engineering, network engineering and storage—each took seriously their individual mandates to provide the best possible solution in their areas of expertise. In those days, there was sometimes friction between the teams when resources held by one team

were needed by another. For example, the systems team might need to tap the storage team for additional LUN space. One did not simply procure one's own resources; there was a carefully considered process (read: red tape) that required diligence, but, eventually, the request would be granted.

Then a strange and wondrous new player came on the IT scene and turned in its head this notion that resources must be managed separately and proprietarily. The rise of virtualization has brought with it the need for a new kind of IT pro—an IT pro with knowledge from all three previously separate areas of IT as well as knowledge of the unique facets of virtualization. One could say that the rise of virtualization has given rise to the one IT pro to rule them all.

**Systems and storage tend to get a lot of attention** in the world of virtualization, but the network remains, as it has always been, the foundational element upon which IT does business. Without the network, the systems could not talk, the printers could not print, and the business could not function. Accordingly, this paper discusses seven networking concepts that the modern virtualization administrator needs to understand in order to maximize the value of the virtualized environment.

**What makes networking in a virtual environment different from networking in the physical world?**

Networks in a virtual environment operate like those in the physical world, but as you start peeling back the virtualization layers, you will find that there has been a shift in how network traffic itself flows around the network. Further, in order to support virtualization-friendly features, administrators have had to make minor changes to the network. This section discusses three key differences between networks in an all-physical world and those in a virtual one.

**Mixed network traffic**

Back in the days of the one-to-one ratio between applications and servers, networking pros had it pretty easy when it came to analyzing network traffic patterns to determine the best design or to troubleshoot a problem. Today, hypervisor-based workload abstraction technology enables IT departments to use hardware more efficiently than ever before, but the technology has added challenges as well. For example, all kinds of mixed network traffic is making its way around the network, making it very difficult for administrators to track down anomalies and plan for traffic peaks and valleys.

In addition, running many virtual machines on a single host has made it a bit more difficult to troubleshoot network problems that arise since the traffic is more difficult to isolate. Further, use of network links in a virtual environment is much greater in one with one application per server.

**Reshaping of the traffic flow**

In addition to adding a lot of mixed traffic to the network, virtualization has fundamentally shifted the way that traffic flows around the network. Whereas network administrators used to see mostly client-to-server communications, which resulted in today's three-tier network architectures, heavy users of hypervisors will see traffic moving not just up and down between client and server, but also side-to-side—that is, hosts communicating with one another.

This traffic pattern shift is caused largely by workload migration technologies, such as vMotion and Live Migration, which administrators use to provide users with a more highly available technology environment. As more hosts are added, more of this side-to-side traffic will occur. This changing of traffic has given rise to a number of new networking technologies, including data center Ethernet (DCE) and software-defined networking (SDN).

In addition to adding a lot of mixed traffic to the network, virtualization has fundamentally shifted the way that traffic flows around the network.



## Stretched VLANs

A long time ago, VLANs were used primarily to enable administrators to break large networks into smaller and more manageable broadcast domains. VLANs were a boon to the network administrator; they could even be considered a very simplistic form of software-defined networking, since VLAN handling is generally handled through the software stack on a network device.

Today, however, VLANs have grown even more important and, as a result, they are being stretched further and further. The reason? vMotion and other workload migration features. These workload migration services are extremely sensitive to latency and operate at Layer 2 of the OSI model (discussed later in this paper). As a result, they are heavily dependent on VLANs being available between all hosts. These hosts could be separated by distances that would not be served in a traditional networking environment.

## How many physical NICs are enough?

One question often asked by fresh virtualization administrators is, "How many physical network interface cards (NICs) are required to support the intended workloads and the services necessary for high availability and workload migration?"

Unfortunately, the answer is a solid "it depends."

It depends on the types of applications that are running in the environment and on the hypervisor features that the organization has chosen to deploy. It also depends on the level of redundancy desired by the organization.

For starters, it's generally recommended that you have at least two network adapters for front-end communication with virtual machines. By bonding multiple network adapters together, effective network capacity is multiplied

while redundancy is added to the environment.

In addition, you may need to consider adding one or more network adapters to support services like vMotion, Fault Tolerance and virtual machine management.

However, with I/O virtualization technologies such as Xsigo available on the market, even the existence of physical network adapters is changing, since these tools enable administrators to slice and dice high-capacity connections into smaller chunks and assign these smaller units to individual services.

## What different virtual network adapters are available and which should be used?

vSphere provides administrators with three choices when it comes to network adapter selection:

- **E1000**—The E1000 network adapter is an emulated version of Intel's 82545EM gigabit Ethernet NIC. Most reasonably modern operating systems have driver support for this adapter. This includes Linux kernel versions 2.4.19 and higher, Windows XP Pro x64 and higher, and Windows Server 2003 and later.
- **VMXNET2**—VMXNET2 is known as a paravirtualized network adapter and has no physical counterpart. It's a software-based network adapter available on a subset of guest operating systems, including Windows Server 2003, Windows XP, Red Hat 5, SUSE Linux Enterprise Server 10, and 64-bit editions of Ubuntu. VMXNET2 provides advanced performance-improving networking features, including Jumbo Frame support and support for hardware offloads.
- **VMXNET3**—VMXNET3 is a brand new paravirtualized network adapter that boasts all of the features of VMXNET2, but adds many more, including Receive Side Scaling, IPv6 offloads, and MSI/MSI-X interrupt delivery. Further, VMXNET3 appears to Windows to be a 10 Gb network adapter, which can have performance advantages, even on slower networks. As long as an

The number of physical NICs you need depends on the types of applications you run, the hypervisor features you have chosen to deploy, and the level of redundancy you desire.

When possible, you should consider the use of the VMXNET3 network adapter in your virtual machines, since VMWare's testing has shown it performs much better than VMXNET2 while consuming less CPU cycles.

organization is running a recent version of vSphere, most reasonably modern guest operating systems are well supported by VMXNET3, including Windows XP and later, Windows Server 2013 and later, and a number of popular Linux variants.

When possible, you should consider the use of the VMXNET3 network adapter in your virtual machines, since VMWare's testing has shown it performs much better than VMXNET2 while consuming less CPU cycles.

### What basic networking concepts do VM admins need to understand and why?

Virtualization has abstracted workloads away from the hardware and has changed the network paradigm. However, in order to adequately support the virtualization needs of the organization, the virtualization administrator should understand some basic networking concepts.

#### Layers 1, 2 & 3 of the OSI networking model

First of all, virtualization administrators should understand the basic concepts behind how the first three layers of the OSI networking model operate and know what kinds of devices operate at each level of the model. With this understanding, administrators can much more easily troubleshoot network communications problems that arise in the virtual environment.

Here's some guidance to assist in this understanding:

- **Layer 1: the physical layer**—Layer 1 is where the most basic communication takes place on the network. It is also where the network cabling and any devices that directly repeat network communications, such as Ethernet hubs, exist. An Ethernet hub has no intelligence whatsoever; its sole purpose is to listen for network traffic on one port and then simply repeat that traffic out all other ports.
- **Layer 2: the data link layer**—This layer is responsible for packetizing Ethernet traffic along with a MAC address and sending

it out on the network. From a device perspective, this is the layer at which switches operate. Further, this is the layer at which VLANs operate. Layer 2 traffic is not capable of leaving the local network or VLAN unless some kind of Layer 3 device is added to the network.

- **Layer 3: the network layer**—This layer is where advanced protocols such as TCP/IP are overlaid atop the Layer 2 network in order to more logically split up a network and control traffic flow. Routers operate at Layer 3.

#### Subnetting

A virtualization administrator must be able to, at the very least, identify when an IP address is local to a vSphere host or a virtual machine. Otherwise, the virtualization administrator can't begin to troubleshoot network connectivity issues or understand how traffic may flow from one virtual machine to another or between vSphere hosts.

A subnet is defined through a combination of an IP address assigned to a resource and the subnet mask assigned to that resource. By performing what is known as a "bitwise AND" using the two variables, a network administrator can tell whether traffic will stay local or whether it needs to be sent out through the network's default gateway for delivery to a computer on a remote network. Bear in mind that certain hypervisor functions, such as vMotion, won't operate if the traffic has to leave what vSphere considers to be the "local" network.

If you'd like to learn everything there is to know about IP subnetting, try the great tutorials, "IP Addressing and Subnetting."

### What role do VLANs play in a virtual environment?

VLANs are often used to segment networks.

As has become the case in the physical world, VLANs can play as important a part or as insignificant a part as desired by the organization. That said, most

organizations rely heavily on VLANs to segment networks in a variety of different ways. Some choose to segment the network by service type; some choose to segment the network by device type; and still others choose to segment the network by the type of user that will reside on a particular segment.

While all of this segmentation could be handled by physically restructuring the network each time a change need to be made, VLANs were introduced to enable network administrators to virtually “tag” specific network ports and uplinks to other devices with VLAN information in order to avoid the physical effort and expense the task would otherwise take.

#### **VLANs have become important for improving security and network performance.**

Over time, VLANs have become important for improving the overall security of an environment and for reducing the size of a network’s broadcast domain in order to improve the overall performance of the network.

In a virtual environment, the purpose of VLAN’s remains exactly the same as in the physical world. Therefore, virtualization administrators need to understand how VLANs work both on their networks as well as in their hypervisor of choice. Further, administrators need to understand how some hypervisor services, such as vMotion, operate on the network—and VLANs and Layer 2 become critically important here. vMotion is a Layer 2 service that requires that workloads be migrated only between hosts on the same VLAN. If the administrator has properly configured networking—that is, ensured that the host servers have a connected network or VLAN dedicated to vMotion traffic—this won’t be a problem, even if the virtual machines themselves are assigned to completely different VLANs.

#### **How are VLANs extended into the virtual environment?**

There are a couple of different ways that VLANs can be extended into the virtual environment. First, an administrator can choose to mimic the physical network as much as possible by dedicating a single VLAN to a single network adapter in the host server. For each additional VLAN that needs to be added, a new physical network adapter would also need to be added to the host.

Obviously, this method won’t scale very far. Therefore, the more common way to extend VLANs into the environment is through the use of 802.1q-based VLAN tagging and trunking VLANs into vSphere; then administrators can use vCenter to simply set the VLAN ID on each individual virtual machine, which provides significant flexibility to the environment.

#### **What are some of the differences between the standard virtual switch and the vSphere Distributed Switch?**

##### **Standard virtual switches**

The whole point of virtualization is to simplify the IT environment and reduce costs through more effective use of hardware. However, the virtual environment itself sometimes adds administrative complexity to the technology management equation.

Take virtual switches (vSwitches) in vSphere, for instance. While a necessary part of the environment, vSwitches must be individually configured on each vSphere host in the environment. For a solution that is supposed to bring some ease to the systems management world, this is quite a bit of effort.

That said, the freely available vSwitch actually does quite a lot, including:

- Forwarding of Layer 2 frames
- Traffic segmentation using VLANs
- 802.1q VLAN tagging
- Support for NIC teaming
- Outbound traffic shaping

VLANs have become important for improving the overall security of an environment and for reducing the size of a network’s broadcast domain in order to improve the overall performance of the network.



A vDS is a cluster-level object that is shared across all of the hosts in the cluster and is managed as a single distributed entity with vCenter. This significantly reduces the administrative networking burden and further abstracts workloads from the host.

### The new vSphere distributed switch (vDS)

The vSwitch pretty well covers the basics, but with vSphere 4, VMware introduced a new type of switch. Available only in the Enterprise Plus edition of vSphere, the new vSphere Distributed Switch (vDS) can be used to further simplify network administration, and it also adds some additional networking capabilities.

Unlike a vSwitch, a vDS is not defined at the host level. Instead, the vDS is a cluster-level object that is shared across all of the hosts in the cluster and is managed as a single distributed entity with vCenter. This significantly reduces the administrative networking burden and further abstracts workloads from the host.

In addition, a vDS:

- Can shape inbound traffic
- Provides support for private VLANs
- Adds support for Netflow to provide visibility into inter-virtual machine traffic
- Adds the ability to mirror a port

But, again, bear in mind that the vDS is available only in the Enterprise Plus edition of vSphere.

### Summary

Today's virtualization administrators need a wide breadth of expertise, including not only virtualization technologies but also systems and storage. Still, networking remains the foundation of the IT infrastructure. By understanding the seven key networking concepts presented in this paper, you can help ensure that your virtual environment operates at peak efficiency and at the lowest cost possible.

### About the Author

Scott D. Lowe is the Founder and Managing Consultant of The 1610 Group, a strategic and tactical IT consulting firm based in the Midwest. Scott has been in the IT field for close to twenty years and spent ten of those years in filling the CIO role for various organizations. Scott is also a micro-analyst for Wikibon and an InformationWeek Analytics contributor. In addition, Scott has also written thousands of articles and blog postings, and he regularly contributes to such sites as TechRepublic, Wikibon, and virtualizationadmin.com. He has also either authored or co-authored four books and is the creator of ten video training courses for TrainSignal.

## For More Information

© 2013 Dell, Inc. ALL RIGHTS RESERVED. This document contains proprietary information protected by copyright. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose without the written permission of Dell, Inc. ("Dell").

Dell, Dell Software, the Dell Software logo and products—as identified in this document—are registered trademarks of Dell, Inc. in the U.S.A. and/or other countries. All other trademarks and registered trademarks are property of their respective owners.

The information in this document is provided in connection with Dell products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Dell products. EXCEPT AS SET FORTH IN DELL'S TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT,

DELL ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL DELL BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF DELL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Dell makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Dell does not make any commitment to update the information contained in this document.

## About Dell

Dell Inc. (NASDAQ: DELL) listens to customers and delivers worldwide innovative technology, business solutions and services they trust and value. For more information, visit [www.dell.com](http://www.dell.com).

If you have any questions regarding your potential use of this material, contact:

## Dell Software

5 Polaris Way  
Aliso Viejo, CA 92656  
[www.dell.com](http://www.dell.com)

Refer to our Web site for regional and international office information.

