



Windows Server 2012: Networking

Module 3: Hyper-V Network Virtualization.

Module Manual

Author: James Hamilton-Adams, Content Master

Published: 4th September 2012

Information in this document, including URLs and other Internet Web site references, are subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2012 Microsoft Corporation. All rights reserved.

Microsoft is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries.

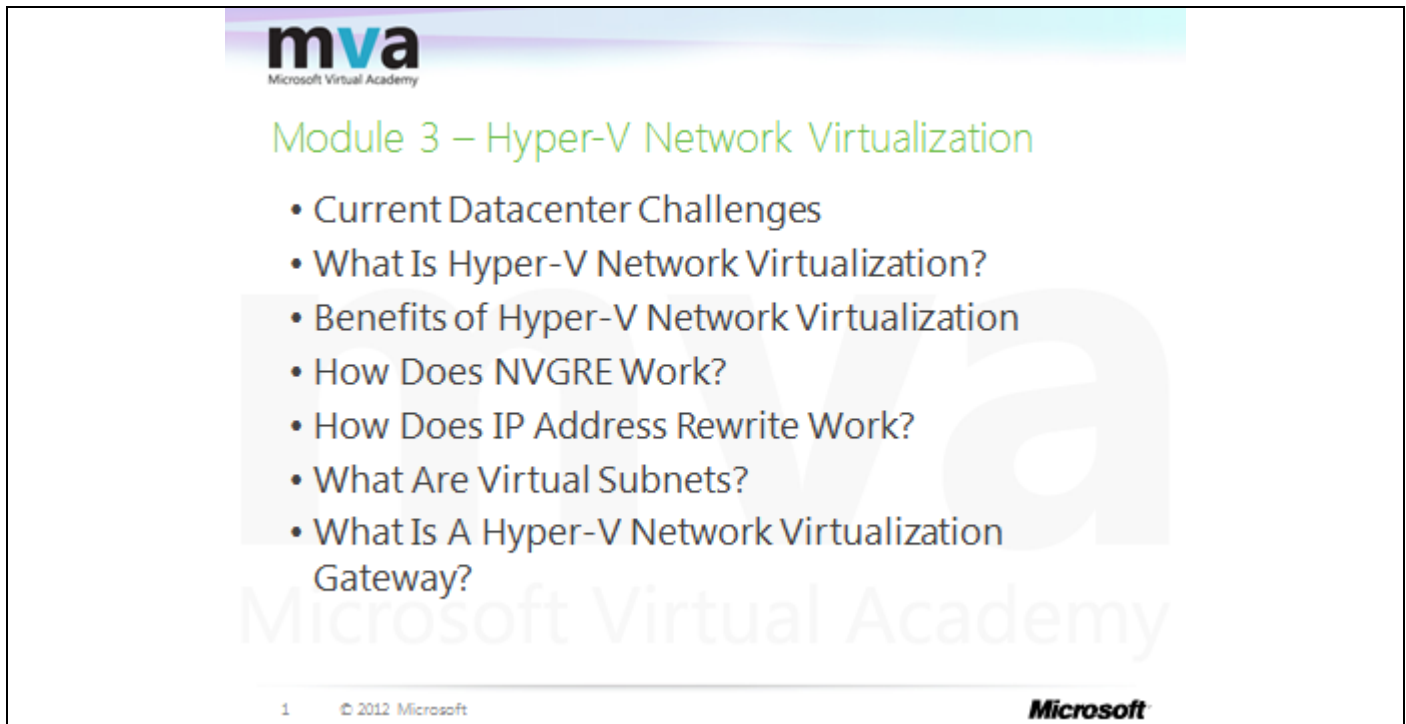
The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Contents

CONTENTS.....	III
MODULE 3: HYPER-V NETWORK VIRTUALIZATION.....	4
Module Overview.....	4
CURRENT DATACENTER CHALLENGES.....	5
WHAT IS HYPER-V NETWORK VIRTUALIZATION?.....	7
BENEFITS OF HYPER-V NETWORK VIRTUALIZATION.....	8
HOW DOES NVGRE WORK?.....	10
HOW DOES IP ADDRESS REWRITE WORK?.....	12
WHAT ARE VIRTUAL SUBNETS?.....	13
WHAT IS A HYPER-V NETWORK VIRTUALIZATION GATEWAY?.....	14
FURTHER READING AND RESOURCES.....	15
<i>Send Us Your Feedback About This Document.....</i>	<i>Error! Bookmark not defined.</i>

Module 3: Hyper-V Network Virtualization.

Module Overview



The slide features the MVA logo at the top left, which includes the text 'mva' in a stylized font and 'Microsoft Virtual Academy' below it. The title 'Module 3 – Hyper-V Network Virtualization' is centered in a green font. A list of seven topics is presented in a blue font, each preceded by a bullet point. The background of the slide is white with a large, faint watermark of the MVA logo. At the bottom left, there is a small number '1' and the text '© 2012 Microsoft'. At the bottom right, the Microsoft logo is displayed.

- Current Datacenter Challenges
- What Is Hyper-V Network Virtualization?
- Benefits of Hyper-V Network Virtualization
- How Does NVGRE Work?
- How Does IP Address Rewrite Work?
- What Are Virtual Subnets?
- What Is A Hyper-V Network Virtualization Gateway?

Hyper-V Network Virtualization is a software-defined networking (SDN) solution included in Windows Server® 2012 that can give the ability to run any service on any server across any cloud. Hyper-V Network Virtualization provides “virtual networks” to virtual machines (VMs), similar to how hypervisor-based server virtualization provides VMs to the operating system. Network virtualization decouples virtual networks from the physical network infrastructure and removes the constraints of virtual LAN (VLAN) and hierarchical IP address assignment from VM provisioning. This flexibility makes it easy for customers to move to infrastructure as a service (IaaS) clouds and efficient for hosters and datacenter administrators to manage their infrastructure, while maintaining the necessary multi-tenant isolation and security requirements and supporting overlapping VM IP addresses.

This virtualization of the network layers promises to enable simpler and automated network configuration, dynamic tenancy isolation, federation of private and public networks, and live migration of VMs among many other benefits.

Microsoft® System Center 2012 SP1 adds support for SDN and Hyper-V Network Virtualization. System Center 2012 SP1 adds support for isolated tenant networks, IP virtualization, switch extensions, and logical switches. Also, with Windows Azure™, Windows Server 2012, and System Center 2012 SP1, you can build an infrastructure that spans private, public, and hybrid clouds.

Current Datacenter Challenges

mva
Microsoft Virtual Academy

Current Datacenter Challenges

- Workload Mobility**
 - Physical location affects IP address
 - Cannot move VM to any host
- Resource Utilization**
 - VM placement not entirely flexible
 - Overprovision hardware to accommodate limitations
- Operational Inefficiency**
 - Coordination needed between network and server teams
 - Increases time to deploy
- Scalable Multi-Tenancy**
 - VLAN limitations not well suited to cloud-based infrastructure
 - Switch reconfigurations on live network add risk
- Onboarding**
 - Organizations moving to a cloud need to reconfigure VMs
 - IP addressing tied to access and security

2 © 2012 Microsoft **Microsoft**

In current server datacenter environments, there are a number of challenges that can prevent administrators from achieving optimum flexibility in a wide variety of scenarios.

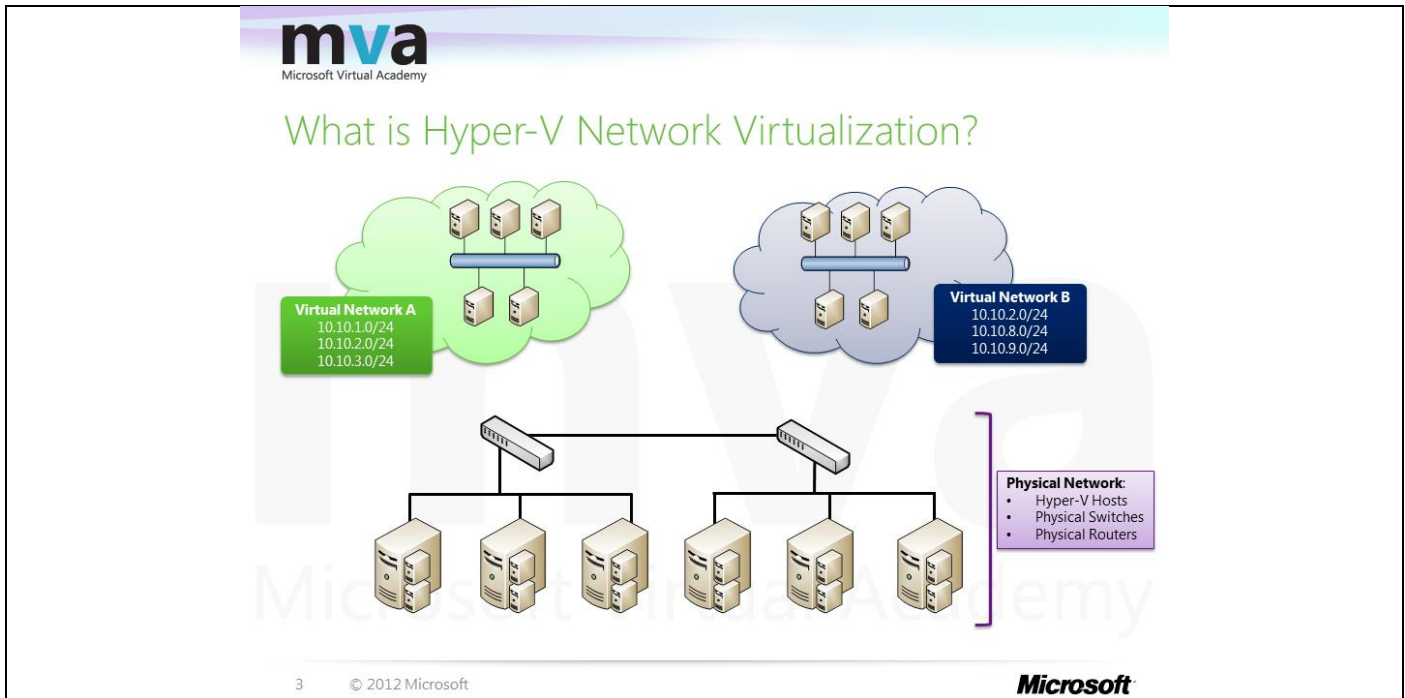
The most common challenges include:

- **Workload mobility.** At present, the physical location directly affects the IP address, so your IP addressing scheme limits VM placement flexibility. You cannot move a VM to a new subnet or physical location without reconfiguring the IP addressing within the VM operating system.
- **Resource utilization.** As a result of the limitations in VM placement and mobility that are dependent on physical location, you do not have complete flexibility to adjust workloads to your VM host servers. This typically results in overprovisioning of hardware resources to ensure that workload requirements are met.
- **Operational inefficiency.** Deploying VMs typically requires tight cooperation between the server team and the network team, but this level of cooperation reduces agility due to the coordination and planning required between these teams.
- **Scalable multi-tenancy.** Typically, VLANs are not well suited for the dynamic requirements of modern private, public, or hybrid cloud environments. Switch reconfiguration in live environments carries an increased risk of network issues or outages.
- **Onboarding.** Organizations that consider moving some or all of their VM infrastructure to a hosted offering currently have to consider reconfiguration of all IP addressing for the hosting environment, and many organizations use current IP address information for access and security rules or policies.

Module 3: Hyper-V Network Virtualization

All of these issues are directly related to the IP addresses assigned to VMs in the network and affect private and hosting datacenters.

What Is Hyper-V Network Virtualization?



Hyper-V Network Virtualization is an SDN solution that enables you to tackle flexibility challenges in your datacenters by abstracting VM IP addresses from the underlying physical network addressing and routing, similar to the way in which VMs are abstracted from the hardware of the physical host they are running on.

In abstracting the IP addressing of VMs, Hyper-V Network Virtualization provides the following advantages:

- You can run multiple virtual networks on one physical network, without needing to use VLANs.
- You can use overlapping IP address ranges in separate virtual networks, even reusing the same IP addresses.
- VM IP addressing is no longer dependent on the location of the VM.
- VMs and their applications do not need to be aware of the network virtualization.
- Each virtual network appears to be a physical network to the VMs that are connected.

With Hyper-V Network Virtualization, each virtual network adapter is associated with two IP addresses:

- **The customer address.** The customer address is the address assigned to the network adapter within the VM guest operating system.
- **The provider address.** The provider address is the address assigned to a physical adapter on the Hyper-V host computer. The provider address is not visible to the VM.

Hyper-V Network Virtualization stores each customer address in a mapping table to send traffic to the correct physical host, which can then forward the traffic to the correct VM.

Benefits of Hyper-V Network Virtualization

The slide features the Microsoft Virtual Academy (mva) logo at the top left. The title 'Benefits of Hyper-V Network Virtualization' is centered in green. Below the title are four colored boxes, each representing a stakeholder group and their benefits:

- Workload Owners (Green box):**
 - Seamless cloud migration
 - Move topology to cloud
 - Preserve policies, settings, IP addressing
 - Simplify test to production migration
- Enterprises (Blue box):**
 - Datacenter consolidation
 - Hybrid cloud capability
 - Easily integrate new acquisitions
- Hosting Organizations (Purple box):**
 - Bring your own IP
 - Bring your own topology
 - Scalable multi-tenancy
- Datacenter Administrators (Dark Blue box):**
 - Flexible VM placement
 - Server and network team agility

At the bottom left of the slide, there is a small number '4' and the text '© 2012 Microsoft'. The Microsoft logo is at the bottom right.

Using Hyper-V Network Virtualization across your VM infrastructure provides the following benefits for the workload owners:

- You can perform seamless migration to cloud hosting for VMs, and VMs do not need reconfiguration before migrating to the cloud.
- You can move many VMs and their associated network topology to the cloud. All addressing information and subnet membership can be preserved.
- You can preserve existing IP dependent policies, and access controls do not require reconfiguration.
- You can reuse IP addresses for test networks or systems and simplify moving VMs between test and production networks.

For enterprises, Hyper-V Network Virtualization provides the following benefits:

- You can consolidate datacenters and still achieve full performance requirements; there is no further need to overprovision hardware.
- You can make use of cloud hosting for some VMs while other VMs reside on-premises and still communicate. This combines the public and private cloud into a hybrid cloud.
- VMs can be moved to another environment, such as a private or public cloud, without reconfiguration. This helps in acquisition and server consolidation scenarios, in addition to organizations moving VMs to hosting environments.

For hosting organizations, Hyper-V Network Virtualization provides the following benefits:

- You can offer “bring your own IP addresses” for VM hosting, enabling customers to move to the public cloud more easily.
- Customers can move entire subnets to the public cloud and still retain the original network topology.

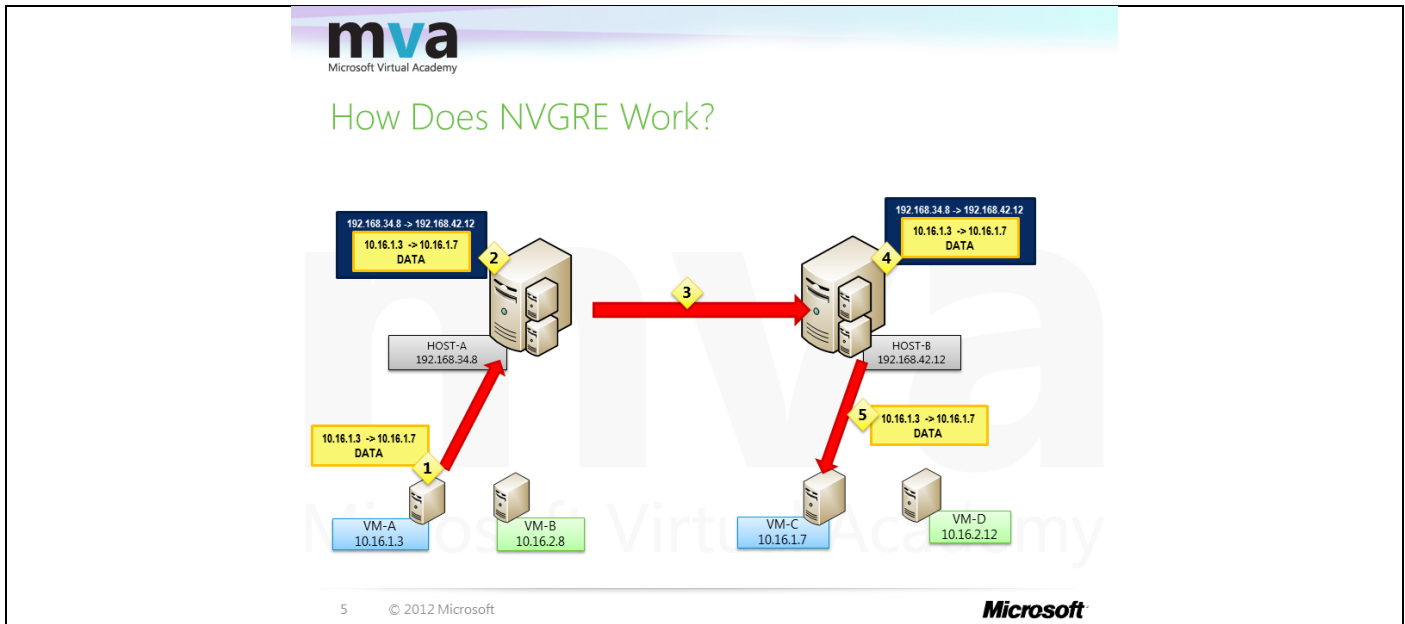
Module 3: Hyper-V Network Virtualization

- Larger datacenters can scale without VLAN limitations.

For datacenter administrators, Hyper-V Network Virtualization provides the following benefits:

- Flexible VM placement on hosts in any physical location. Moving VMs does not require reconfiguration.
- Server and network teams can work more independently to improve agility.

How Does NVGRE Work?



Hyper-V Network Virtualization can use two different mechanisms to virtualize IP addresses:

- Network Virtualization Generic Routing Encapsulation (NVGRE).
- IP address rewriting.

NVGRE encapsulates (or tunnels) IP packets from VMs inside another packet before transmission on the physical network. The encapsulation packet has physical network IP addressing and routing information that corresponds to the Hyper-V hosts on the physical network.

In the diagram, VM-A has an IP address of 10.16.1.3 and runs on HOST-A; HOST-A has an IP address of 192.168.34.8 on the physical network. VM-C has an IP address of 10.16.1.7 and runs on HOST-B, which has an IP address of 192.168.42.12 on the physical network. When VM-A sends information to VM-C, the following sequence occurs:

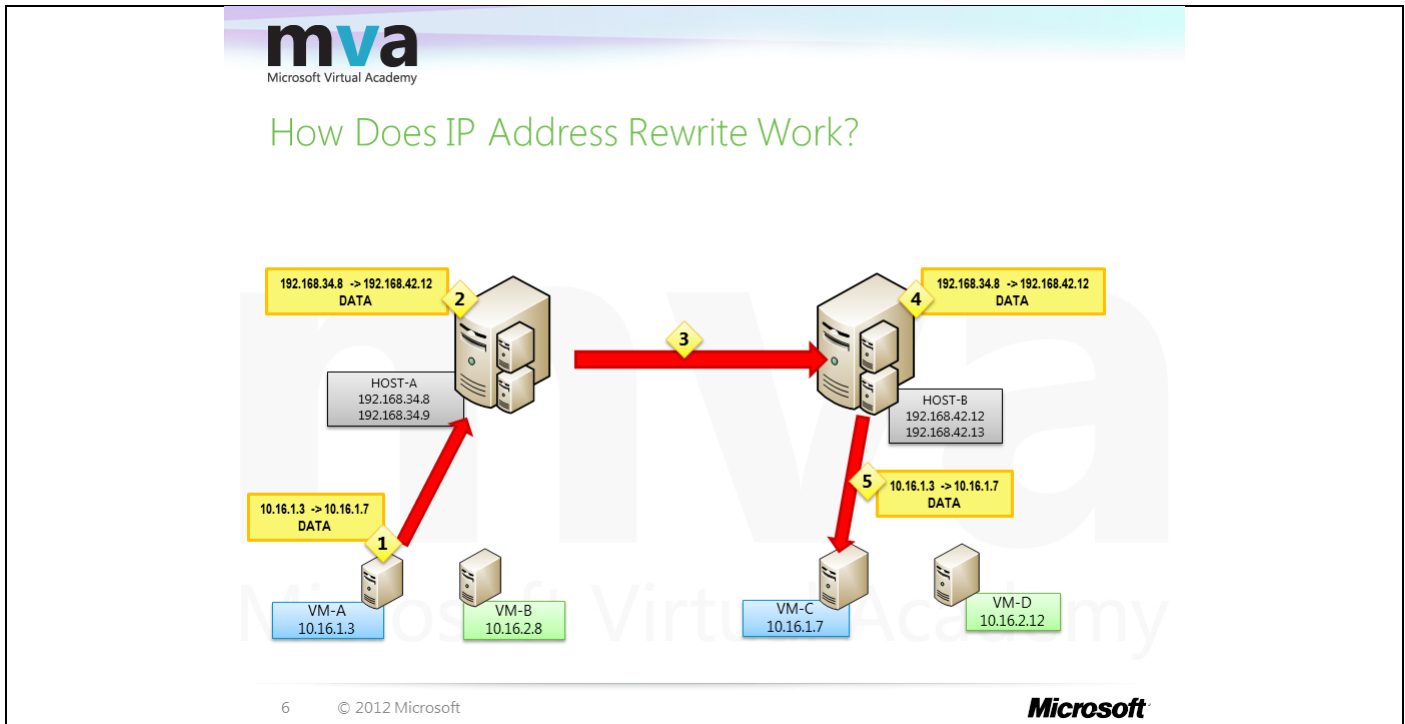
1. VM-A sends a TCP packet to the network with a source IP address of 10.16.1.3 and a destination address of 10.16.1.7.
2. The TCP packet arrives at the Hyper-V network switch on HOST-A. Due to a Hyper-V Network Virtualization policy applied to HOST-A, the TCP packet is encapsulated in an NVGRE packet. The new NVGRE packet has a source address of 192.168.34.8 and a destination address of 192.168.42.12. The NVGRE packet also has a value representing the virtual subnet ID (VSID) added to it.
3. The NVGRE packet is sent over the physical network and arrives at HOST-B.
4. HOST-B examines the VSID and unwraps the NVGRE packet. Based on the VSID and the destination address of the original TCP packet, HOST-B then forwards the TCP packet to the Hyper-V virtual switch on HOST-B connected to VM-C.
5. The TCP packet arrives at VM-C.

VM-A and VM-C are never aware of the NVGRE addressing or the VSID. VSIDs can enable secure multi-tenant hosting, where different tenants can reuse the same IP address ranges, but within their own virtual subnet.

Module 3: Hyper-V Network Virtualization

NVGRE does not require any special hardware; you can deploy NVGRE using your existing switch and routing infrastructure. NVGRE is the preferred network virtualization mode for most deployments.

How Does IP Address Rewrite Work?



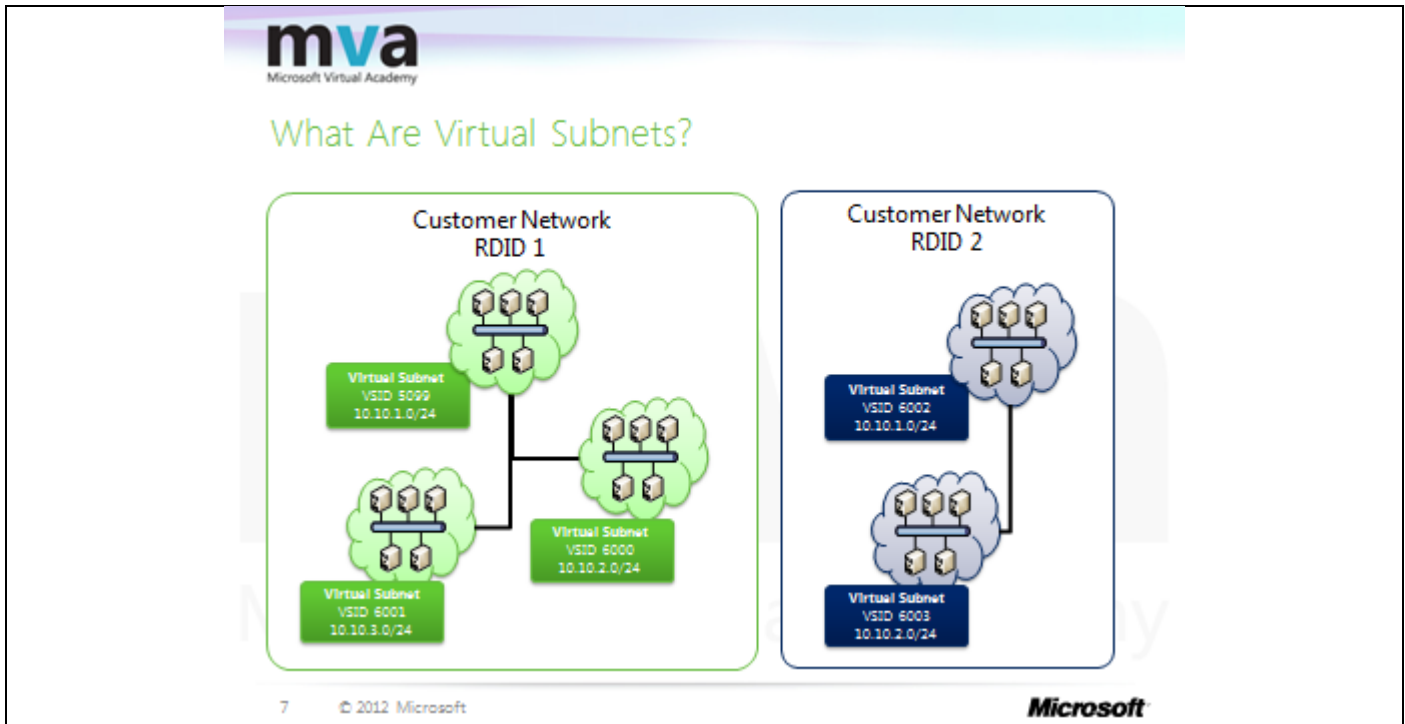
You can also use the IP rewrite mechanism for network virtualization. IP address rewrite changes the source and destination addresses in a TCP packet when transmitting data between VMs over the physical network.

In the diagram, VM-A has an IP address of 10.16.1.3 and runs on HOST-A; HOST-A has IP addresses of 192.168.34.8 and 192.168.34.9 on the physical network. VM-C has an IP address of 10.16.1.7 and runs on HOST-B, which has IP addresses of 192.168.42.12 and 192.168.42.13 on the physical network. When VM-A sends information to VM-D, the following sequence occurs:

1. VM-A sends a TCP packet to the network with a source IP address of 10.16.1.3 and a destination address of 10.16.1.7.
2. The TCP packet arrives at the Hyper-V network switch on HOST-A. Due to a Hyper-V Network Virtualization policy applied to HOST-A, the TCP packet has the source address changed to 192.168.34.8 and the destination address changed to 192.168.42.12.
3. The TCP packet is sent over the physical network and arrives at HOST-B.
4. HOST-B examines the TCP packet and, based on the destination address, HOST-B then forwards the TCP packet to the Hyper-V virtual switch on HOST-B connected to VM-C.
5. The TCP packet arrives at VM-C.

IP address rewrite is intended for VM workloads with very high bandwidth requirements, such as 10 gigabits per second. IP address rewrite mode requires a unique provider address for each customer address.

What Are Virtual Subnets?



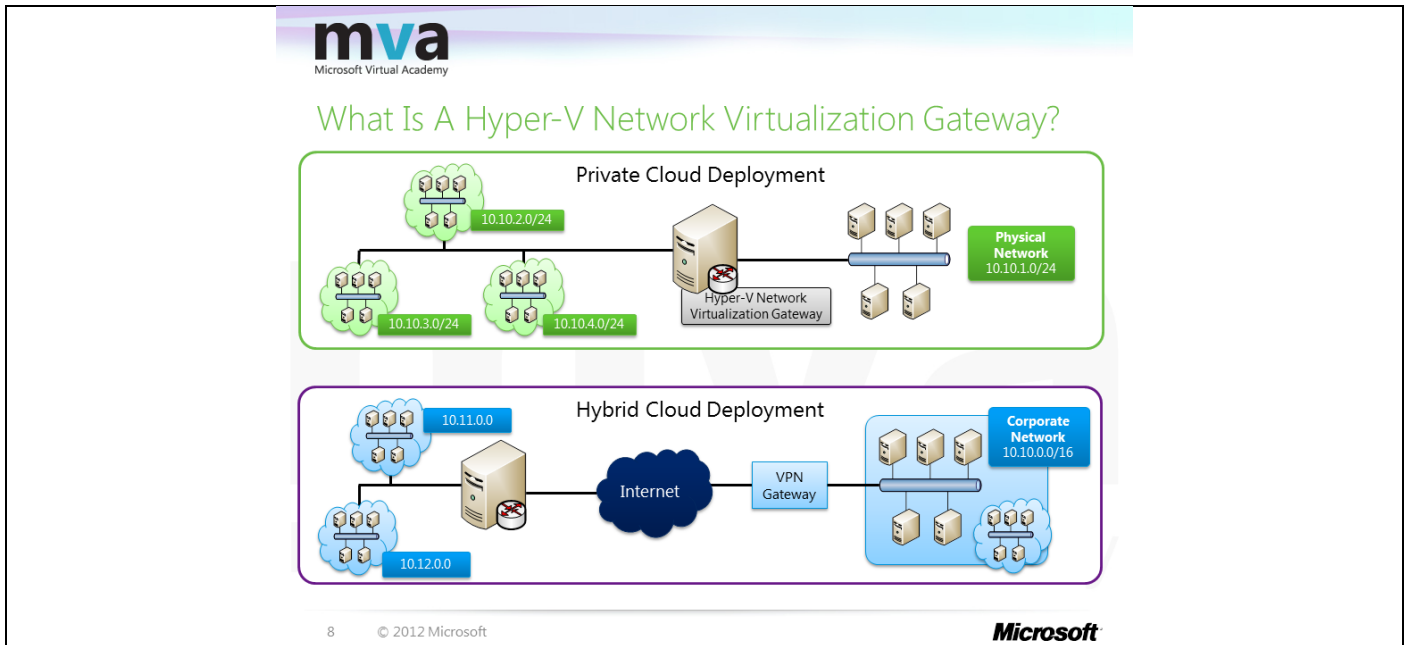
In Hyper-V Network Virtualization, you can create virtual subnets to fulfill similar requirements to a VLAN in physical networking. A virtual subnet functions as a broadcast domain, segregating IP traffic from other virtual subnets. Virtual subnets provide the following behavior:

- VMs in the same virtual subnet must use the same IP prefix, although virtual subnets accommodate both IPv4 and IPv6 addresses simultaneously.
- Each virtual subnet belongs to a single entity called a customer network and is assigned a single VSID. The VSID is a unique value between 4096 and $(2^{24}-2)$.

In Hyper-V Network Virtualization, because we use a customer/provider model to describe different elements of the virtual networking, we can refer to a customer network that represents a number of virtual subnets. A customer network is identified by a unique routing domain ID (RDID). A customer network forms an isolation boundary because Hyper-V Network Virtualization will not deliver traffic between RDIDs. However, one customer network may have many virtual subnets, and network virtualization can route traffic between virtual subnets with the same RDID.

In the diagram, traffic from the virtual subnet with the VSID 5099 can be delivered to VMs on the other two virtual subnets within the customer network RDID 1. However, traffic from any virtual subnet in customer network RDID 1 cannot pass to virtual subnets in customer network RDID 2.

What Is A Hyper-V Network Virtualization Gateway?



You cannot virtualize every server and client in a network infrastructure, so physical servers, client computers, and network devices must still communicate with VMs on the network. You can use a Hyper-V Network Virtualization gateway to enable VMs on virtual subnets to communicate with hosts in subnets and VLANs on the physical network.

A significant benefit of Hyper-V Network Virtualization is the ability to use virtual private network (VPN) capabilities to connect two different physical locations and pass traffic between VMs in each location over the VPN without the VMs being aware of the VPN. This benefit enables seamless multi-datacenter deployments over the Internet, enabling hosting organizations to provide additional capacity to private organizations when needed, by moving VMs to the hosting infrastructure without reconfiguration. This scenario is called a hybrid cloud deployment, and with Hyper-V Network Virtualization, hosting organizations can allow "bring your own IP address," even if customer addresses overlap.

Further Reading and Resources

[Hyper-V Network Virtualization Overview](#)
[Network Virtualization technical details](#)