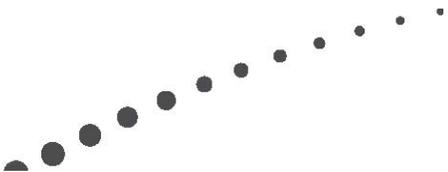




***Citrix XenServer Design:
Designing XenServer Network
Configurations***

www.citrix.com





Contents

About.....	5
Audience.....	5
Purpose of the Guide	6
Finding Configuration Instructions.....	6
Visual Legend	7
Additional Terminology	8
Chapter 1: Introduction	9
Chapter 2: Basic XenServer Networking Concepts.....	11
Introduction to XenServer Networking	11
Connecting Virtual Machines to Networks.....	12
Networking Configuration after Installation.....	14
Impact of Pools on XenServer Networking	15
Sequence of Networking Configuration Tasks.....	18
Cabling Configuration for XenServer	18
Connecting XenServer to Physical Switches	21
Chapter 3: Sample Networking Scenario	22
Example: Adding Virtual Machines to a Network.....	22
Creating Network Resiliency through Bonds.....	23
Connecting a VM to a Network using Virtual Network Interfaces.....	25
Segregating VM Traffic from Management and Storage Traffic.....	27
Scenario 1: Segregating Traffic.....	28
Scenario 2: Using the Management Network for VM Traffic.....	29
Scenario 3: Isolating VM Traffic on a Private Network.....	30
Scenario 4: Connecting VMs to Multiple Linked VLANs	34



Chapter 4: Specifying Networking Requirements.....	40
Overview	40
Introduction.....	41
XenServer Networking Support and Requirements	42
Defining Your Networking Requirements.....	43
Considering Workload Communication Requirements.....	43
Evaluating Your Current Network Configuration	45
Determining Host Networking Requirements.....	46
Reviewing Initial Pool Design against Networking Requirements	48
Calculating the Number of Physical NICs per Host	49
Calculating Bandwidth Requirements	50
Chapter 5: Designing XenServer Networks.....	52
Overview	52
Deciding to Use the Distributed Virtual Switch.....	53
Designing Network Redundancy.....	56
Considering NIC Bonding.....	56
Active-Active Bonding	57
Active-Passive Bonding.....	60
LACP Link Aggregation Control Protocol Bonding	61
Bonding Managed Interfaces and MAC Addressing.....	65
Best Practices for Bonded Interfaces	66
Designing Networks for Performance.....	67
Resolving Network Bottlenecks.....	69
Testing XenServer Network Performance	70
Limiting Bandwidth Consumption for High Demand Workloads.....	72
Additional Considerations	73



Creating a Dedicated Network for Storage Motion.....	74
Enabling Switch-port Locking	74
Enabling Promiscuous Mode for Traffic Monitoring.....	74
Chapter 6: Designing Your Storage Network Configuration.....	76
Overview	76
Connecting to Storage Controllers	77
Creating a Separate Storage Network.....	77
Assigning IP Addresses to NICs (Secondary Interfaces)	79
Configuring Redundancy for Storage Traffic	82
Choosing to Enable Multipathing Support or Bond NICs.....	83
Suggestions for Improving Storage Network Performance	84
iSCSI Storage.....	84
Configuring Networks with Jumbo Frames	85
Chapter 7: Considering Network Performance for PVS – XenServer Deployments	88
Virtualizing the Provisioning Services Server	88
IP Addressing Requirements.....	90
Isolating the Provisioning Services Streaming Service	90
Disabling the Spanning Tree Protocol.....	90
Best Practice Configurations	90
Chapter 8: Verifying Your XenServer Networking Configuration	92
Verifying XenServer Networking on a Pool	92
Verifying your Physical Configuration	92
Verifying your XenServer Networking Settings and Configuration	94
Resolving Issues	95
Revision History.....	97

About

This guide helps you understand design your XenServer networking and design a networking configuration for XenServer environments. It includes the following topics:

- Best practice information about the management interface, NIC bonding, jumbo frames, and storage networks
- High-level information about features you may want to enable as part of your networking configuration, such as the Distributed Virtual Switch solution
- The correct sequence in which to configure XenServer networking, including guidance about cabling XenServer hosts and connecting them to physical switches
- Checklists to help you gather requirements for your XenServer networking configuration

Audience

Before reading this guide, you should have a basic knowledge of networking. This guide has several audiences:

- **Systems Architects.** Systems architects who are designing a virtualized environment.
- **Infrastructure Engineers and Network Administrators.** Networking and storage professionals who configure storage or manage the Layer 2 network infrastructure in their organizations.
- **Application Administrators.** XenApp and XenDesktop administrators who are implementing a virtualization solution to virtualize Citrix products, IT infrastructure, or other applications they manage.



This guide assumes that you are familiar with basic XenServer concepts, including XenServer installation, XenCenter, resource pools, and the pool master.

Purpose of the Guide

This guide is meant to provide you with the best-practice information you need to design your XenServer networks.

To provide you with the foundation you need to understand the recommendations, the first half of the guide provides an explanation of XenServer networking concepts using a scenario-based approach.

The second half of the guide provides you with information to help you select between various XenServer networking options and information about the best ways to configure them.

Because this is a design guide, it generally does not provide configuration instructions except as needed to clarify concepts. As the most common way of managing XenServer and XenServer pools is through XenCenter, this guide mainly refers to XenCenter and XenCenter help, unless specified differently.

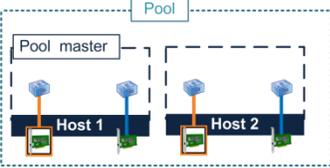
Finding Configuration Instructions

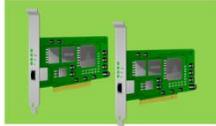
You can find networking configuration instructions in the following locations:

- **XenCenter Help.** The XenCenter help provides UI-based step-by-step instructions using XenCenter, the XenServer UI-based administration console. Users who are not comfortable with the XenServer `xe` commands, may prefer this option.
- **XenServer Administrator's Guide.** The *XenServer Administrator's Guide* provides command-line based instructions for performing networking tasks. For integrators, it also provides information about XenServer networking from the object-model perspective.

Visual Legend

This guide relies heavily on diagrams to explain key concepts. These diagrams use the following icons:

Icon	Meaning
	<p>Virtual Machine. A virtual computer that runs on the XenServer host.</p>
	<p>Virtual Network Interface. On a VM, there is a logical interface that appears and functions like a NIC; this interface is known as a <i>virtual network interface</i>. A virtual network interface lets VMs send and receive network traffic. Some product literature refers to virtual network interfaces as <i>VIFs</i> and <i>virtual NICs</i>.</p>
	<p>Network. A network is the virtual network switching fabric built into XenServer that lets you network your virtual machines. It links the physical NICs to the virtual network interfaces and connects the virtual network interfaces together.</p>
	<p>Host. A XenServer host is the physical server on which the XenServer hypervisor is running.</p>
	<p>NIC. The physical network card (NIC) in your host.</p>
	<p>Pool. A XenServer resource pool or “pool” is a connected group of hosts which provides a platform on which virtual machines run.</p> <p>To join hosts to a pool, they require broadly compatible hardware and must be running the same XenServer version and patches.</p> <p>Pools comprise a pool master and subordinate servers known as pool</p>

	<p>members (sometimes also referred to as "slaves"). The pool master provides a single point of contact for all the servers in the pool and the master will forward commands to individual pool members as necessary.</p>
	<p>Physical Switch. The device on a physical network that connects network segments together.</p> <p>This guide may present physical switches either as a three-dimensional physical box or as a one-dimensional panel with ports.</p>
	<p>NIC Bond. In this guide, enclosing NICs in green represents a bond.</p> <p>A NIC bond is a pair of NICs configured so they logically function as one network card. NIC bonding is also known as <i>NIC teaming</i>.</p>

Additional Terminology

These terms appear in the sections that follow:

Management Interface. The management interface is a NIC assigned an IP address that XenServer uses for its management network, including, but not limited to, traffic between hosts, between a host and Workload Balancing and for live migration.

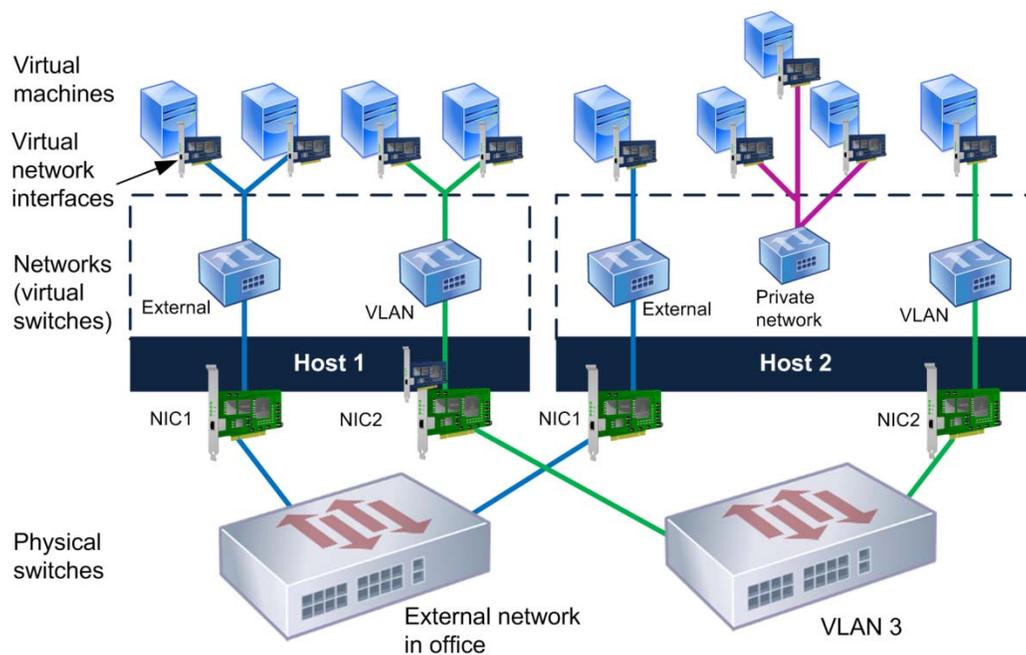
VM traffic. VM traffic refers to network traffic that originates or terminates from a virtual machine. This is sometimes referred to as guest traffic or VM/guest traffic.

Note: For consistency with XenCenter terminology changes in the XenServer 6.1 release, this guide now uses the term *management interface* to refer to the IP-enabled NIC that carries the management traffic. In the last release, this guide used the term primary management interface for this purpose. Likewise, this guide now uses the term *secondary interface* to refer to an IP-enabled NIC configured for storage traffic. Throughout this guide, the term *managed interface* is used to refer to management interfaces and secondary interfaces collectively.

Chapter 1: Introduction

This documentation explains basic networking concepts and their application by using a series of scenarios to illustrate the concepts. The scenarios begin immediately after installation and end with connecting a VM to a network.

These sample scenarios focus on four different types of networks: External Networks, VLANs, cross-server private networks, and single-server private networks. If you configured the scenarios demonstrated in this guide, by the time you finished, you would create a deployment that looked like the following illustration.



This illustration shows how virtual machines connect to three different types of networks: an external network, a VLAN network, and a single-server private network.



This guide explains these types of networks by providing the following information:

Chapter 2 introduces XenServer networking and explains how to prepare for XenServer networking configuration by configuring the physical infrastructure and hardware layers in your environment, including the correct sequence for physically configuring networking. The chapter also discusses the effect pooling XenServer hosts has on networking and describes the networking configuration after installation.

If you want to read a list of XenServer networking definitions before reading this information, see the Visual Legend on page 7. Otherwise, the definitions are provided in chapter 2 as you read the sections.

Chapter 3 provides several sample scenarios that illustrate how to add virtual machines to a network. The first scenario guides you through the process of segregating different types of traffic, including storage and management traffic. The second scenario gives you an alternative to dedicating NICs to specific types of traffic; it shows an example of using the management network for management and VM traffic. The third scenario shows an example of how to segregate traffic by creating a single-server private network on a host.

Chapter 4 provides guidance about how to determine what networking configurations and hardware your XenServer deployment will require. The chapter includes tables with references to relevant information to help you jump to key information about how to define your NIC configuration, calculate bandwidth requirements, and other topics.

Chapter 5 provides information about key design choices you will make when creating your XenServer networks, including whether to use the Distributed Virtual Switch, how to configure XenServer network redundancy (NIC bonding), and how to design XenServer networks for performance.

Chapter 6 provides information about how to create a separate physical network for storage traffic, set an IP address on a NIC, configure iSCSI multipathing, and configure support for jumbo frames. This chapter also includes suggestions for improving storage network performance and best practices.

Chapter 7 provides information about how to virtualize a Provisioning Services server on XenServer using SR-IOV. This chapter also provides some references to best practices for XenServer-Provisioning Services deployments.

Chapter 8 provides information about how to verify your XenServer networking configuration after you physically configured it. This chapter provides a process for verifying networking on a host and on pools as well as information about resolving issues.

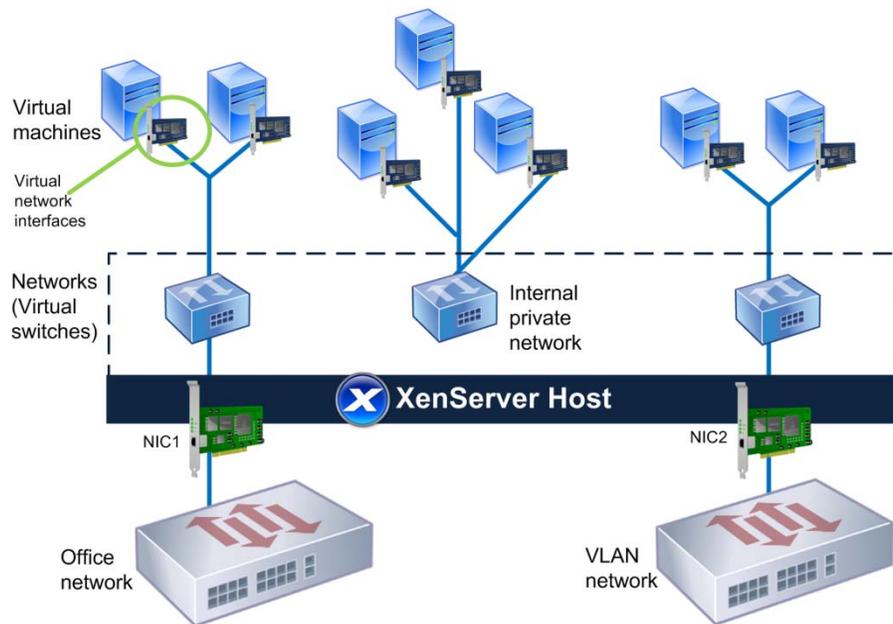
Chapter 2: Basic XenServer Networking Concepts

This chapter includes the following topics:

- An introduction to XenServer networking
- The network settings created during installation

Introduction to XenServer Networking

XenServer provides virtual networking features that let you build networks for your virtual machines the same way you build networks for physical machines.



The VMs connect to three different types of networks: an office network, an internal private network, and a VLAN.

You can connect virtual machines to your production network like you connect physical machines or build private networks within a host or pool for testing, development, or security purposes. You can connect virtual machines to your VLAN networks using standard VLAN configurations.

The most important networking components XenServer lets you configure are *virtual network interfaces* and *networks*:

- **Virtual network interfaces.** Virtual machines connect to networks using virtual NICs, known as virtual network interfaces. Virtual network interfaces let VMs send and receive network traffic. You can assign each virtual network interface its own IP address and MAC address. Some product literature refers to virtual network interfaces as *VIFs* and *virtual NICs*.
- **Networks.** XenServer has an internal virtual switch, known as a network, that lets virtual machines on a XenServer host communicate with each other using the same networking protocols that are used on physical networks.

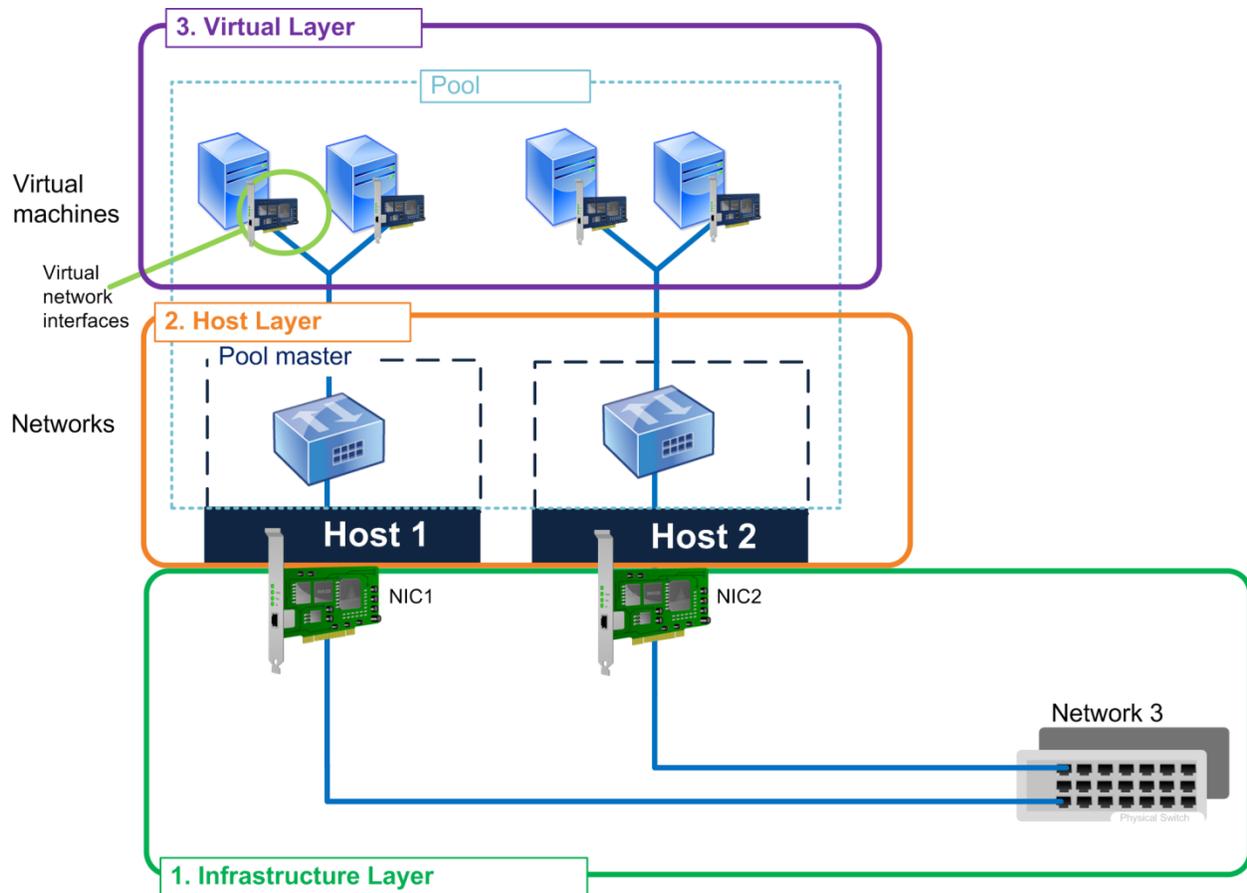
A network is the logical network switching fabric built into XenServer that lets you network your virtual machines. It links the physical NICs to the virtual network interfaces and connects the virtual network interfaces together. These networks are virtual switches that behave as regular L2 learning switches. Some vendors' virtualization products refer to networks as *virtual switches* or *bridges*.

Connecting Virtual Machines to Networks

When you are configuring network connectivity on XenServer hosts, your ultimate goal is to connect the VMs to a network. To do this:

1. Connect the host to a physical network. (For VMs without external network connectivity, you would configure a private network instead.)
2. Connect the VM by creating a Virtual Network Interface for it and connecting the Virtual Network Interface to a network. As shown in the illustration on page 11, the virtual network interfaces on the VMs connect to networks in a host and then connect to a physical network through the host's NIC.

One way to think about these tasks is that you need to configure connectivity at both the hardware and virtual layers as shown in the illustration that follows.



This illustration shows the order in which you should configure networking in your virtual environment: (1) Start on the physical infrastructure layer, which means connecting NICs to switches; (2) configure the hardware layer, which means connecting hosts to networks and configuring these networks; (3) configure the virtual layer, which means attaching VMs to networks through virtual network interfaces.

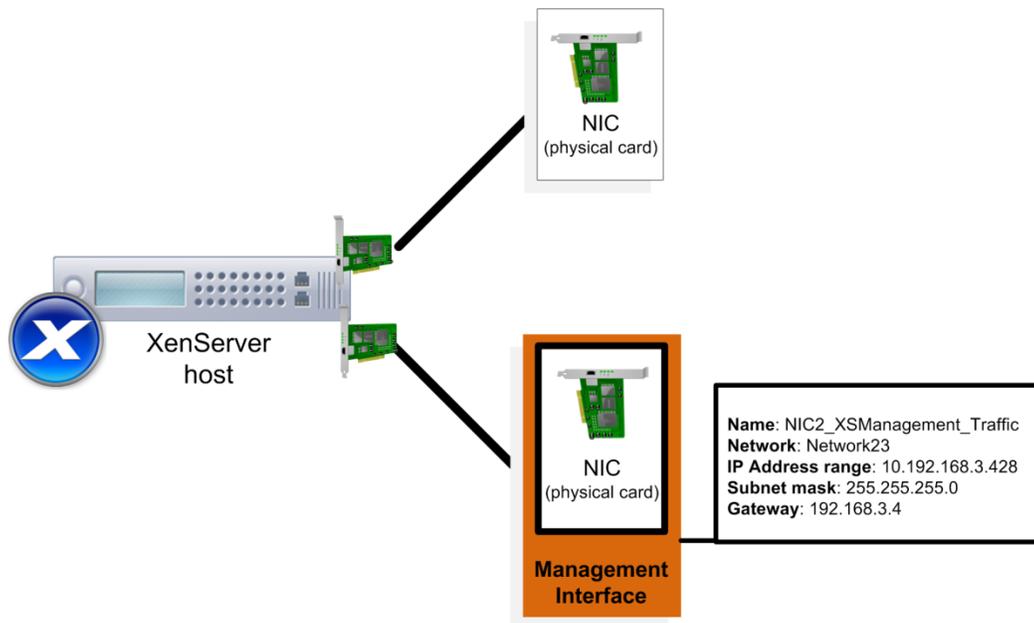
Important: Configuring networking in the order listed described in “Sequence of Networking Configuration Tasks” on page 18 is critical. If you vary from this sequence, the management interface may not be configured correctly on each host. If this occurs, all VMs in the pool may start on the pool master and not their home or optimal servers.

Networking Configuration after Installation

After installation, the XenServer *host* has all the information it needs to connect to at least one of your external networks. This is because you define the following networking options while installing XenServer:

- IP Address Configuration and Other Settings.** You set the host's initial XenServer networking configuration when you first install XenServer on the physical computer. XenServer Setup configures options, such as the IP address configuration (DHCP/static), based on the values you provide during installation.
- Network Connectivity.** XenServer installation prepares each NIC connected to a switch for network connectivity by creating one network for each NIC. This means that if the host has, for example, three NICs, XenServer creates three networks: Network 0, Network 1, Network 2. For a visual explanation, see page 15.
- Management Interface and the Management Network.** During XenServer Setup, you specify an IP address for one NIC. XenServer uses that NIC to connect to your organization's network and to carry management traffic for functions like communicating with other hosts in a pool, XenCenter, Workload Balancing, and other components. This NIC is known as the *management interface*. This is the only NIC that Setup configures with an IP address.

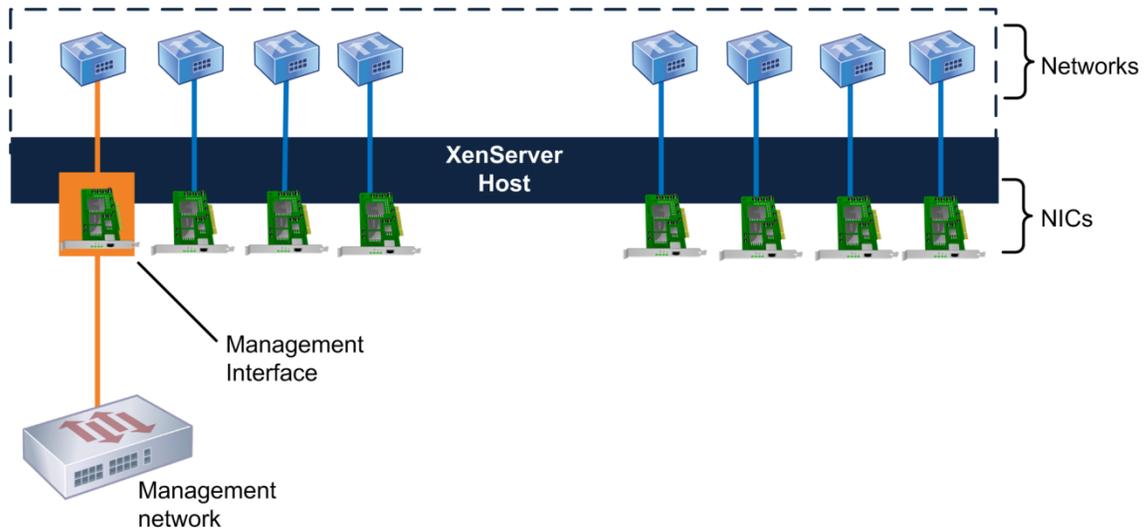
The illustration that follows shows a regular (unconfigured) NIC and a NIC configured as a management interface.



This illustration contrasts a regular NIC with one configured as the management interface. The management interface has an IP address, subnet mask, and gateway assigned to it.

During installation, XenServer also creates a separate network for each NIC it detects on the host. Unless you change this set up, XenServer uses the additional NICs on the host for VM traffic only.

The illustration that follows shows an example of XenServer’s initial network configuration following installation.



This illustration shows how, during installation, XenServer makes you choose a NIC as the management interface. In this case, the administrator selected NIC0. XenServer uses the other NICs for VM traffic.

Most environments require additional configurations to these basic network settings. These can range from creating pools to integrating additional networks, connecting your VMs to those networks, and configuring a separate storage network. The scenarios in the following chapter provide examples of these tasks.

Note: If you plug any NICs into switches after installing XenServer, if you cannot see the NICs in XenCenter or xsconsole, you might need to either a) run `xe pif-list` or `xe pif-plug` in the CLI or reboot the XenServer host.

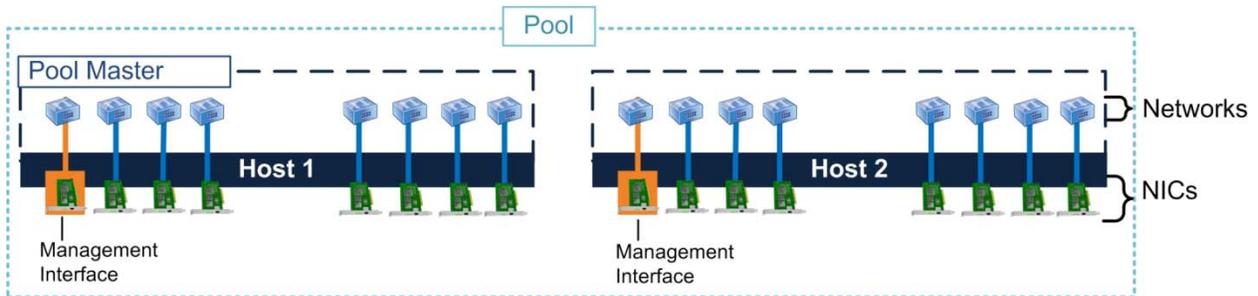
Impact of Pools on XenServer Networking

Networking is a pool-level feature in XenServer. When you change networking on the pool master, XenServer synchronizes all hosts in a pool to use the same network settings.

As a result, for XenServer to operate correctly, you must ensure that network settings match across all hosts in the pool, including:

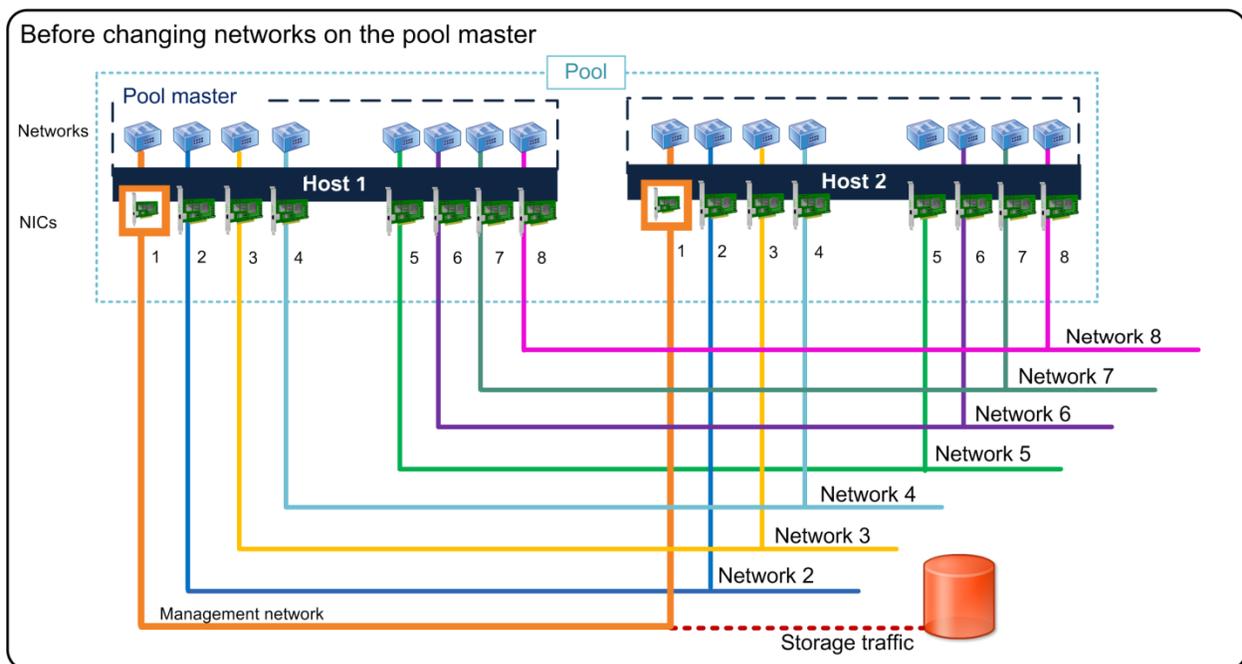
- Which NICs are bonded
- Which NICs are configured as the management interface
- Which NICs connect to storage

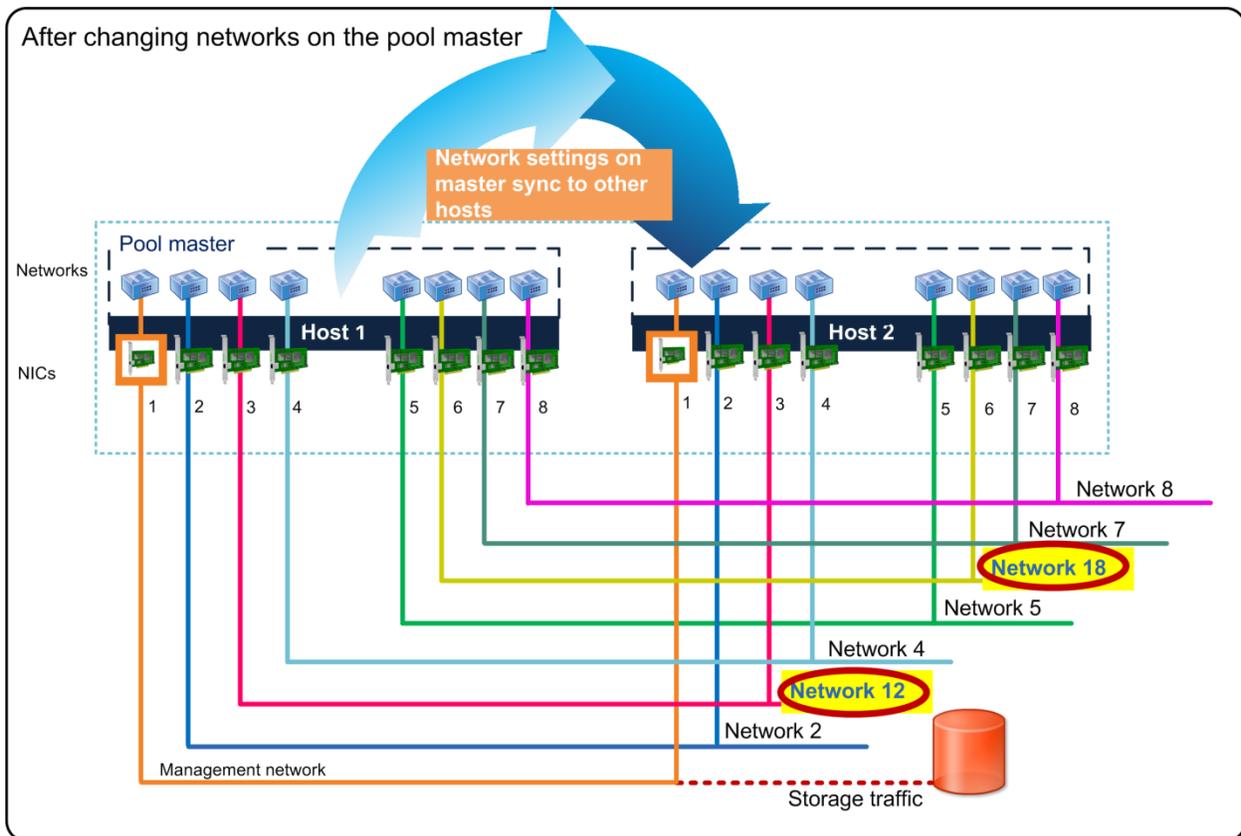
The networks to which NICs connect must be the same on the corresponding NICs on each host in the pool.



This illustration shows two hosts joined together in a pool before any networking configuration is performed on them.

Ideally, you should add all desired hosts to the pool before configuring any network settings. Pooling the hosts before configuring networking creates cleaner records in XenServer’s internal networking-configuration database.





These two illustrations show how XenServer replicates the network settings created on the pool master on all other hosts in the pool. In the top illustration, NICs 3 and 6 on both hosts use Networks 3 and 6. In the bottom illustration, after reconfiguring NIC 3 on the pool master to use Network 12 and NIC 6 to use Network 18, XenCenter automatically configures the other host in the pool to use those settings.

After creating a new pool or joining a host to an existing pool, XenServer automatically replicates the network settings on the master to the joining hosts.

When you use XenCenter to make networking changes, XenCenter changes the other hosts to match the newly modified host. When you use the CLI to change network settings, you must either:

- Change each host manually to match the modified host's settings
- Make the change on the pool master and restart all the member hosts in the pool

XenServer requires network settings to match across the pool because of features that use live migration, such as XenMotion, High Availability, and Workload Balancing. These features enable the physical server hosting a VM to change at any time, and possibly automatically without your intervention. Therefore, the VMs must be able to access all of their target networks regardless of which host XenServer moves them on to.



For this reason, it is critical to have *and maintain* an identical physical cabling, NIC, and switch configuration for each host across the pool. Likewise, Citrix strongly recommends changing the physical configuration on all hosts in a pool before changing network settings on the pool.

Important: After joining the hosts to the pool, check the management interface on each member host to make sure that it has its own unique IP address and/or set the correct static IP address.

Sequence of Networking Configuration Tasks

Citrix recommends performing your initial networking configuration in the sequence that follows to help ensure XenServer stores your networking configuration correctly:

1. Cable the hosts by plugging all NICs into the appropriate switches, as described in “Cabling Configuration for XenServer” on page 18.
2. Configure the switches. See “Connecting XenServer to Physical Switches” on page 21.
3. Install XenServer on the hosts. Citrix recommends that you ensure your networking configuration is set up correctly before creating a resource pool, since it is usually easier to recover from a bad configuration in a non-pooled state. To verify networking is set up correctly, see “Chapter 8: Verifying Your XenServer Networking Configuration.”
4. Create a pool of the hosts, if you want to pool them. See “Impact of Pools on XenServer Networking” on page 15.
5. Configure NIC bonds and networks. For more information, see the scenarios in “Chapter 3: Sample Networking Scenario.”

Important: Do not configure the XenServer High Availability feature until after you complete your networking configurations. Networking configuration can interrupt the High Availability heartbeat and cause hosts to shut themselves down. Consequently, the hosts probably will not reboot correctly and will need the **host-emergencyha-disable** command to recover.

Cabling Configuration for XenServer

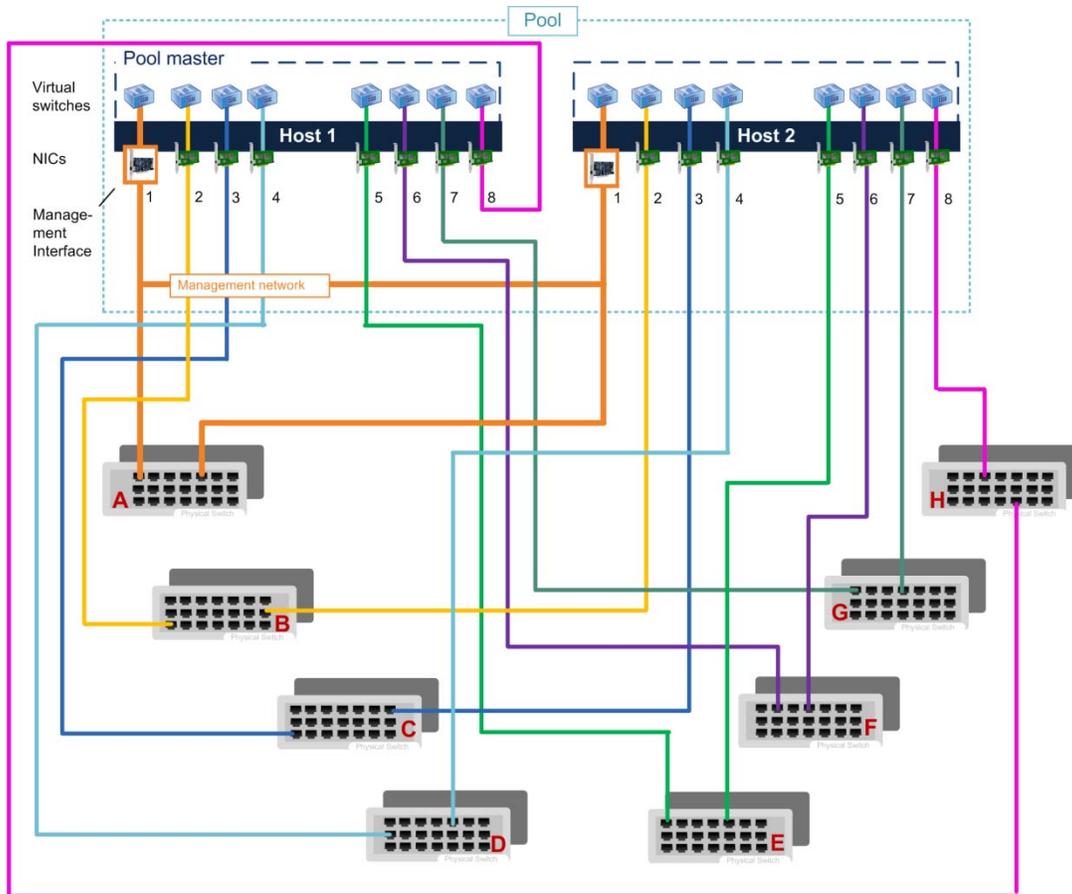
Citrix recommends plugging the physical Ethernet cables into all the NICs and the appropriate switches **before** installing XenServer. The ideal process is as follows:

1. If you did not cable your hosts before installation, plug all the NICs in each host in the pool into the appropriate switch ports.
2. Connect the corresponding NICs on each host in the pool to the same physical switch (that is, the same subnet).

The term *corresponding* refers to the NIC of the same number on another host. For example, NIC 3 on Host 1, NIC 3 on Host 2, NIC 3 on Host 3. This means that each individual NIC

on every host must connect to the same physical network as the NIC in the same position on all other hosts in the pool.

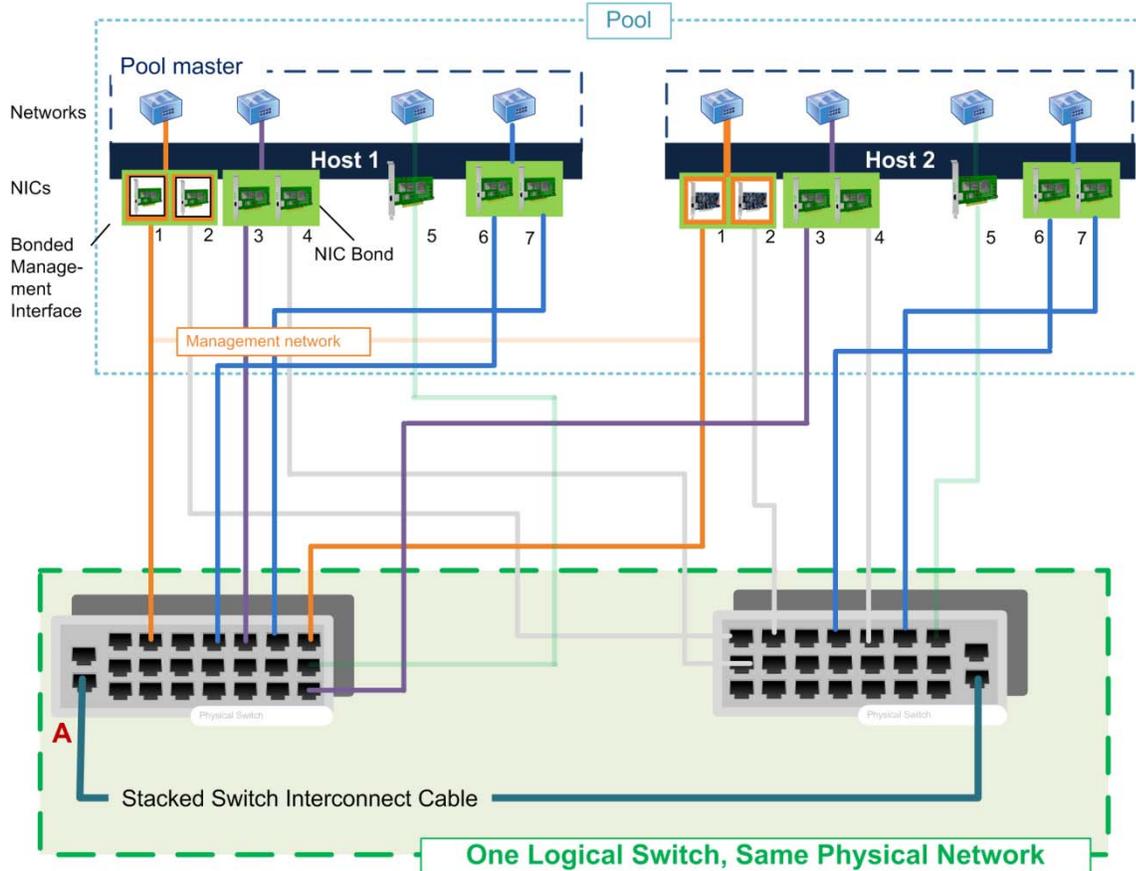
The following figure is a visual example of this configuration in an enterprise environment.



This illustration shows how each corresponding NIC on both hosts must physically connect to the same network. Each switch represents a separate physical network. Each member host's NICs must be connected to the same physical networks as the corresponding NICs on the pool master.

Ensuring the cabling on each host in the pool is correct is critical. As shown in the previous illustration, all NICs must connect to the same physical networks (shown as separate switches) as the NICs in the same position on all hosts across the pool.

In an environment with only one logical switch (for example, one that has a hierarchy of switches that form one large physical network), you only need to connect the NICs to switches on that network that have the same physical or logical (VLAN) connectivity. The example that follows shows how you might cable such an environment.



*This illustration shows two stacked switches connected with the vendors proprietary interconnect cable across the same physical network. These switches function logically as one unit. Because there are no VLANs configured on any of the ports and all ports have the same connectivity, the NICs can be plugged into any port on these two switches. Note: The need for stacked switches varies according to bond mode, as described on page **Error! Bookmark not defined.***

XenServer cannot detect if you make any errors while setting up the physical network. For example, if a XenServer host expects to be able to contact a specific gateway using a certain NIC, XenServer cannot indicate the cabling is incorrect. If you receive errors, they might not indicate network configuration as the cause.

Ensuring that the corresponding NIC on each host has the same network configuration is what ensures that a host's VM attached to, for example, Network 1, can communicate with a VM attached to Network 1 on another host. This ensures that if you migrate a VM to a new host, the VM retains the same physical connectivity after migration.

Note: When you configure networking, if you do not have all of your NICs plugged in to switches, you must have, at a minimum, the NIC(s) for the management interface on all hosts in your pool plugged into your network. Otherwise, the pool master cannot synchronize its network settings to the member hosts. Likewise, if you are using a dedicated NIC for storage, you must also connect the cables for that NIC on each host.



Connecting XenServer to Physical Switches

When you are connecting a XenServer host to a switch, configure the switch's ports differently than you would when connecting a workstation to a switch. There are specific, critical guidelines about the Spanning Tree Protocol (STP) and enabling PortFast.

For more information, see CTX123158—[*Considerations for XenServer Switch Ports*](#).

Chapter 3: Sample Networking Scenario

This chapter provides a scenario-based example of how to connect virtual machines to a physical network. This includes the following:

- Segregating traffic
- Using the management network for traffic in a very small environment

Example: Adding Virtual Machines to a Network

This section provides a sample scenario of a simple networking configuration that includes connecting VMs to networks, creating redundancy, and configuring NICs.

Designing a XenServer networking deployment may require several tasks, including, for example, configuring redundancy for network availability, configuring NICs, and, ultimately, connecting VMs to the desired networks. During this process, you might also separate different types of traffic for security or performance reasons (for example, separating traffic for managing the XenServer platform from VM traffic).

Before configuring networking on a pool, you should know to which networks your VMs will need to connect. A standard network configuration process might require:

1. Configuring redundancy for network availability.
2. Creating separate storage or management networks (used to separate management or storage traffic from VM traffic).
3. Creating VMs and connecting them to the desired XenServer network(s).

This section provides you with an example of that process. This section describes the different configuration options and steps required to put your virtual machines on the network by using a

sample scenario. While the scenario might not directly apply to your environment, it is designed to put XenServer's networking features into context.

Creating Network Resiliency through Bonds

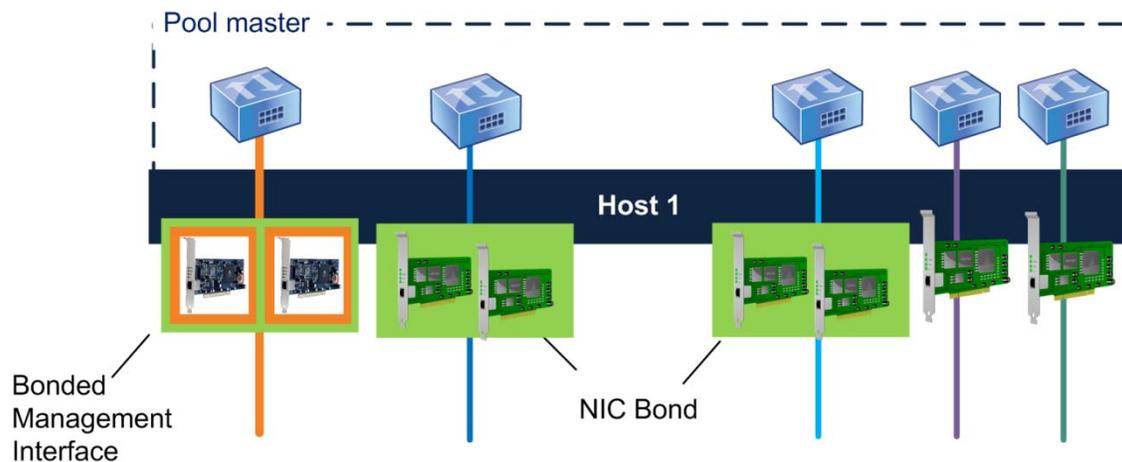
After joining all hosts to your pool, you may want to ensure that any critical servers have high availability access to the network. One way XenServer lets you achieve high network availability is to create redundancy through *NIC bonding*.

NIC bonding is a technique for increasing resiliency and/or bandwidth in which an administrator configures two or more NICs together so they logically function as one network card. Both NICs have the same MAC address and, in the case of management interfaces and secondary interfaces, have one IP address.

XenServer supports *bonding* two or more NICs together on a host. If one NIC in the bond fails, XenServer automatically redirects traffic to the second NIC. NIC bonding is also sometimes known as *NIC teaming*.

You can use XenCenter or the xe CLI to create NIC bonds. If XenCenter is managing a pool, XenServer automatically replicates the bonding configuration across all hosts in the pool.

In the illustration that follows, the management interface is bonded with a NIC so that it forms a bonded pair of NICs. XenServer will use this bond for management traffic.



This illustration shows three pairs of bonded NICs, including the management interface. Excluding the management interface bond, in this example, XenServer uses the other two NIC bonds and the two un-bonded NICs for VM traffic.

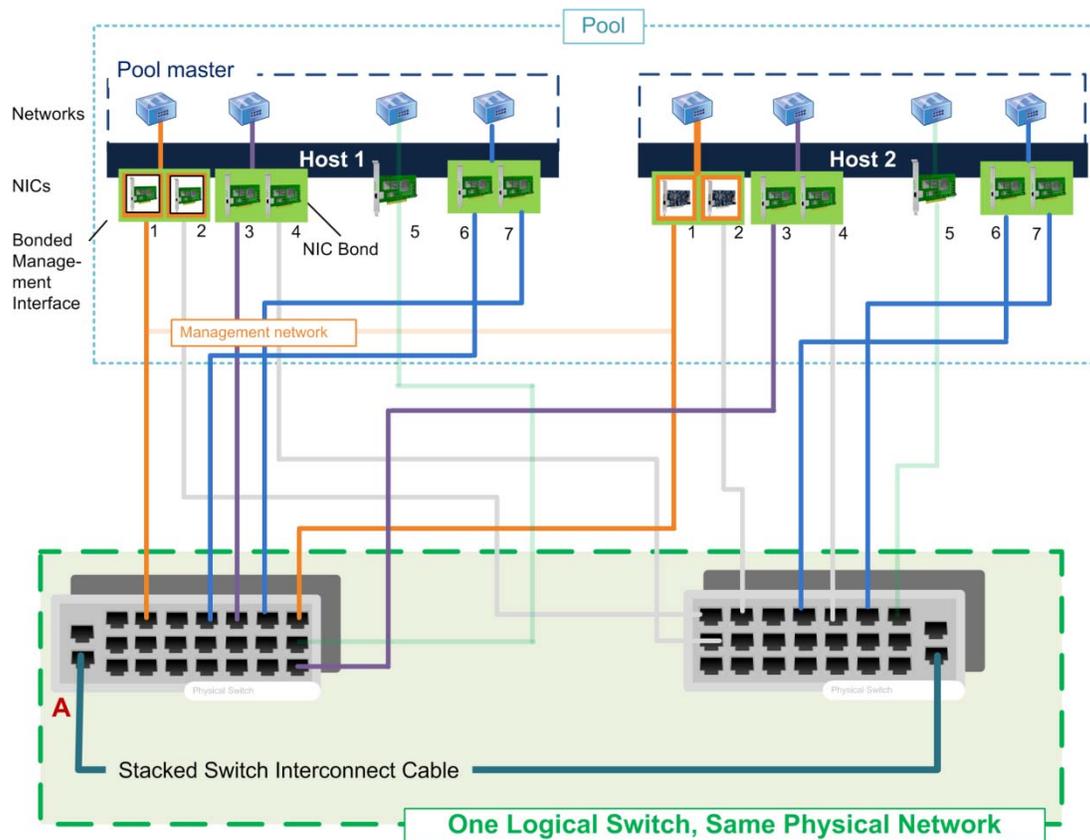
Ensuring Resilience through Redundant Switches

When VM networks use bonded NICs, traffic is sent over both NICs. If you connect one of the NICs in a bond to a second (redundant switch) and a single NIC or switch fails, the virtual machines remain on the network since their traffic fails over to the other NIC/switch.

Provided you enable bonding on NICs carrying only guest traffic, both links are active and NIC bonding can balance each VM's traffic between NICs. Likewise, bonding the management interface NIC to a second NIC also provides resilience. However, only one link (NIC) in the bond is active and the other remains unused unless traffic fails over to it.

If you bond a management interface, a single IP address is assigned to the bond. That is, each NIC does not have its own IP address; XenServer treats the bond as one logical connection.

The illustration that follows shows how the cables and network configuration for the bonded NICs have to match.



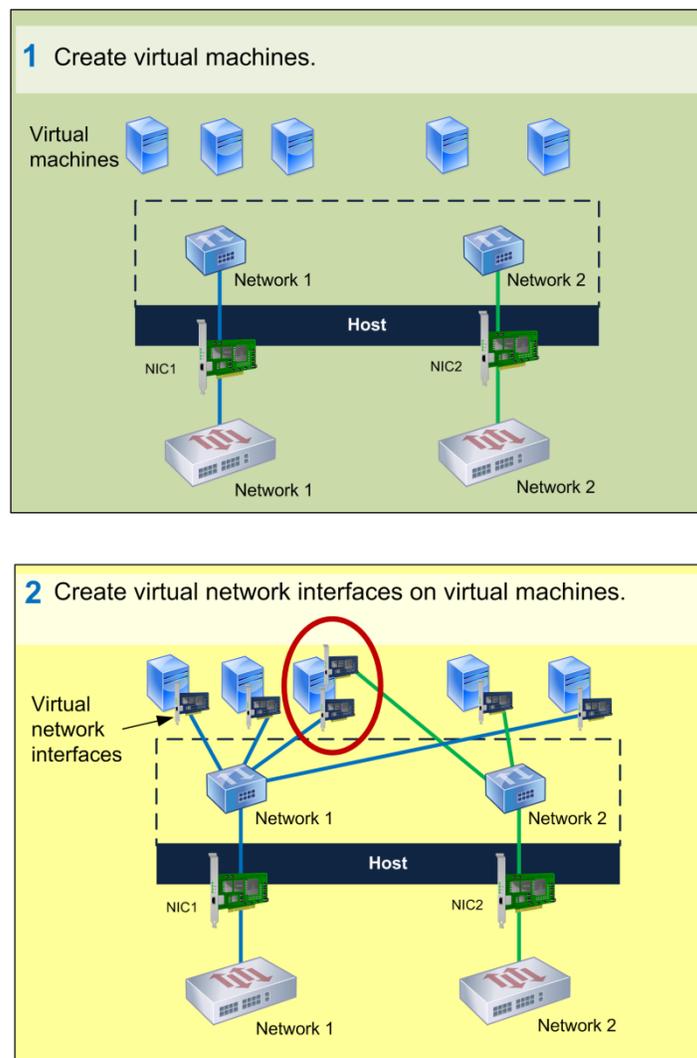
This illustration shows how two NICs in a bonded pair (LACP bond mode) use the same network settings, as represented by the networks in each host. The NICs in the bonds connect to two stacked switches for redundancy.

Note: For more information about bonds, see “Considering NIC Bonding” on page 56.

Connecting a VM to a Network using Virtual Network Interfaces

Virtual machines connect to a network through a virtual network interface on that particular network. XenServer sends the VM's traffic through the target network's associated NIC. By default, when you create a VM in XenCenter, XenServer creates a virtual network interface connecting the VM to Network 0. This configuration lets VMs connect to an external network through the NIC attached to Network 0.

You need a virtual network interface on a VM for each separate physical network to which you want to connect it. In environments that connect to only one physical network, the virtual network interface XenCenter creates by default when you create a VM may be sufficient for your needs. However, if you need a VM to connect to multiple physical networks, you must create a virtual network interface for each one of those networks.



This illustration shows how VMs require a virtual network interface for each physical network to which they need to connect.



Some additional points about virtual network interfaces:

- Most, but not all, VMs have at least one virtual network interface. (If an administrator accesses a VM only through XenCenter, the VM does not need a virtual network interface.)
- Each virtual network interface must have a “virtual” MAC address. You can configure XenServer to generate these automatically for you (recommended) or specify them manually.
- When you create a network in XenCenter, you can specify if you want XenCenter to create a new virtual network interface for that network automatically, whenever you create a VM.
- Unlike for the physical and infrastructure layers, the networking configurations on VMs do not need to match other VMs in the pool.

Note: To determine which VM is associated with a virtual network interface, see CTX122520—[How to Find which Virtual Network Interface is Assigned to a Virtual Machine in XenServer](#).

Understanding Virtual MAC Addressing

Just like NICs in the physical world, each virtual network interface must have its own (virtual) MAC address. When you create a virtual network interface, you can either specify a MAC address manually or let XenServer generate one for you.

When XenServer generates MAC addresses automatically, it generates *locally administered addresses*. Locally administered addresses are addresses assigned to devices by a user, which typically lack manufacturer-specific encoding. As a result, they do not contain a manufacturer-specific *Organizationally Unique Identifier* (OUI). Typically, manufacturers “burn-in” MAC addresses in which the first three octets indicate which company manufactured the device.

This means that the MAC addresses XenServer generates will not clash with addresses from *hardware* devices on your network.

XenServer generates a MAC addresses at random based on the random seed in the *VM.other-config:mac-seed* parameter of the VM and the device number of the virtual network interface (a sequence number for the VIF: 0...6).

A particular combination of a MAC seed and device number always results in the same MAC address. Consequently, if you remove a virtual network interface from a VM and recreate it later, the new virtual network interface typically gets the same MAC as before.

XenServer preserves MAC addresses when migrating VMs. However, when you copy or clone VMs, the VM receives a new random MAC address seed and the virtual network interfaces get new MAC addresses based on that seed.

Tip: To obtain the MAC address of a XenServer VM in XenCenter, select the VM’s **Network** tab, select the virtual network interface, and click **Properties**.

Segregating VM Traffic from Management and Storage Traffic

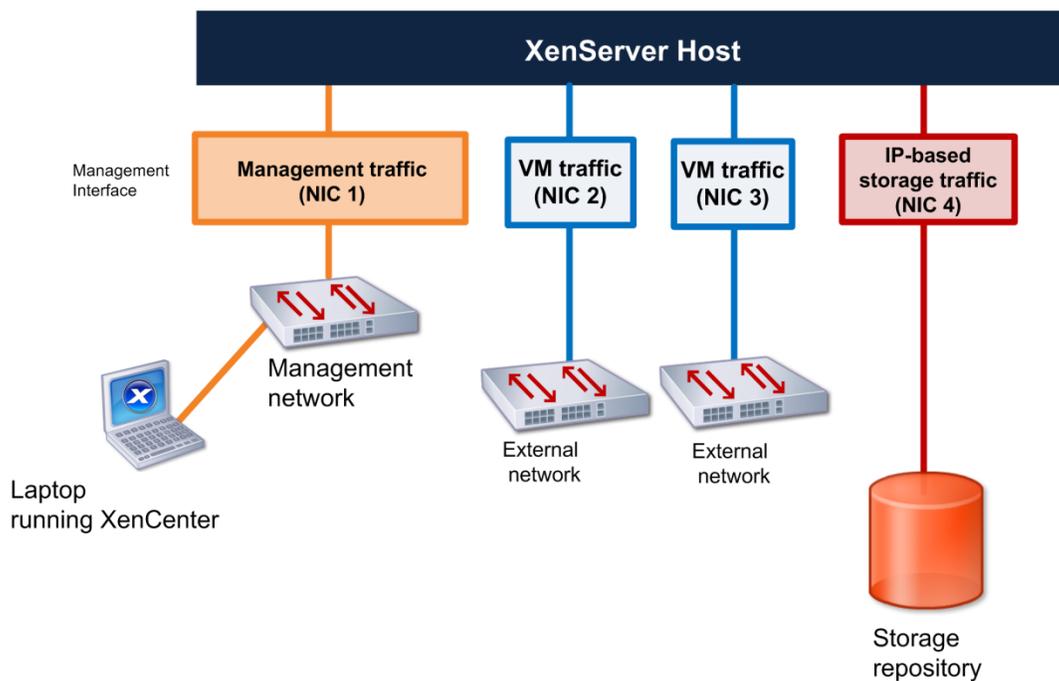
You can separate each type of traffic –VM, storage, and management traffic – onto its own network for either security or performance reasons.

For most environments, Citrix recommends segregating VM traffic from management traffic as the best practice. Not only does it increase the security of the management network, it can improve performance by reducing competition between traffic types for network resources, reducing potential collisions, and reducing the load on the management interface.

There are a variety of ways in which you can separate traffic, including:

- Separating all types of traffic from each other. For example, putting the virtual machines on a network not used for storage or management traffic.
- Separating the management traffic from the VM and storage traffic.

However, VMs will only use a NIC for VM traffic if they have a virtual network interface on the same network as the NIC. The illustration that follows shows the best practice example of how you might separate traffic.



This illustration shows how NICs that are not designated for management or storage traffic only carry VM traffic.

While separating traffic is a best practice in larger environments, it is not an absolute requirement for all environments. In smaller environments, you may want to configure VMs to send their traffic on the management network. However, Citrix recommends evaluating the performance of this configuration regularly.

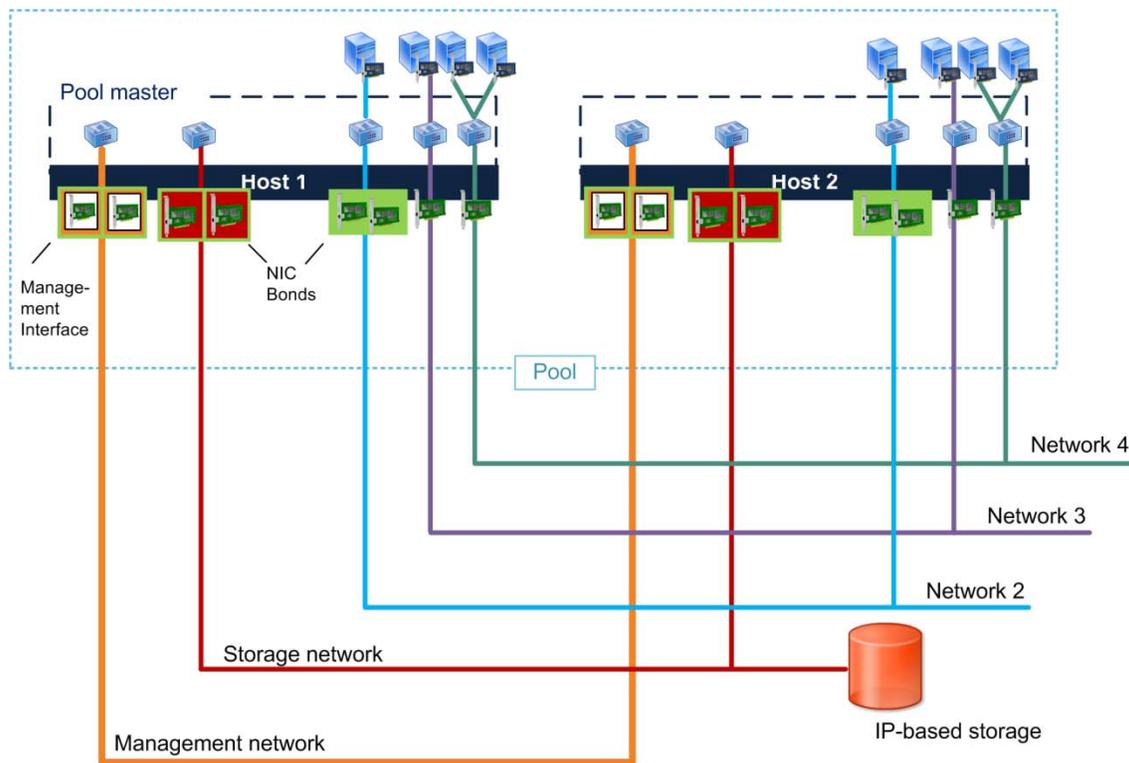
The scenarios that follow illustrate both of these concepts: separating traffic and sending traffic over NICs shared by multiple networks.

Scenario 1: Segregating Traffic

In this scenario, an administrator wants a dedicated network for management and storage traffic. To do this, the administrator:

- Attached the network cables coming from the NICs to a switch for a network to be used for VM traffic, which is physically isolated from the storage and management networks
- Created virtual network interfaces on the same networks as the NICs

The illustration that follows shows these segregated networks.



This logical illustration shows segregated guest, storage, and management networks. In this scenario, all the VMs using network 2 can communicate with each other because they are configured to use the same (corresponding) NIC bond on their respective hosts and that bond connects to the same physical network. Likewise, the two VMs connected to network 3 can communicate with each since the corresponding NIC 7 on each host connects to the same physical switch.

As shown in previous illustration, not all NICs have virtual network interfaces associated with them. If you do not configure a virtual network interface connecting to the management network, the management NIC becomes dedicated for management traffic. For example, in the previous



illustration there are NICs connected to the management and storage networks that do not have corresponding virtual network interfaces.

Note: Citrix does not recommend assigning IP addresses (that is, creating management interfaces) for each NIC on your host. Ideally, Citrix does not recommend using any NICs with IP addresses assigned to them for VM traffic.

Scenario 2: Using the Management Network for VM Traffic

In environments with minimal security requirements, you can configure VMs to share the management or storage networks.

In this example, the organization uses the management network for two purposes:

- XenCenter can connect to the management network through the management interface on the pool master. This is because of the IP address on that NIC. Likewise, hosts and other components, such as Workload Balancing, can use the connection to communicate with XenServer.

Note: XenCenter only communicates with the pool master and not any member servers. Specifically, XenCenter only connects to the IP address of the master's management interface.

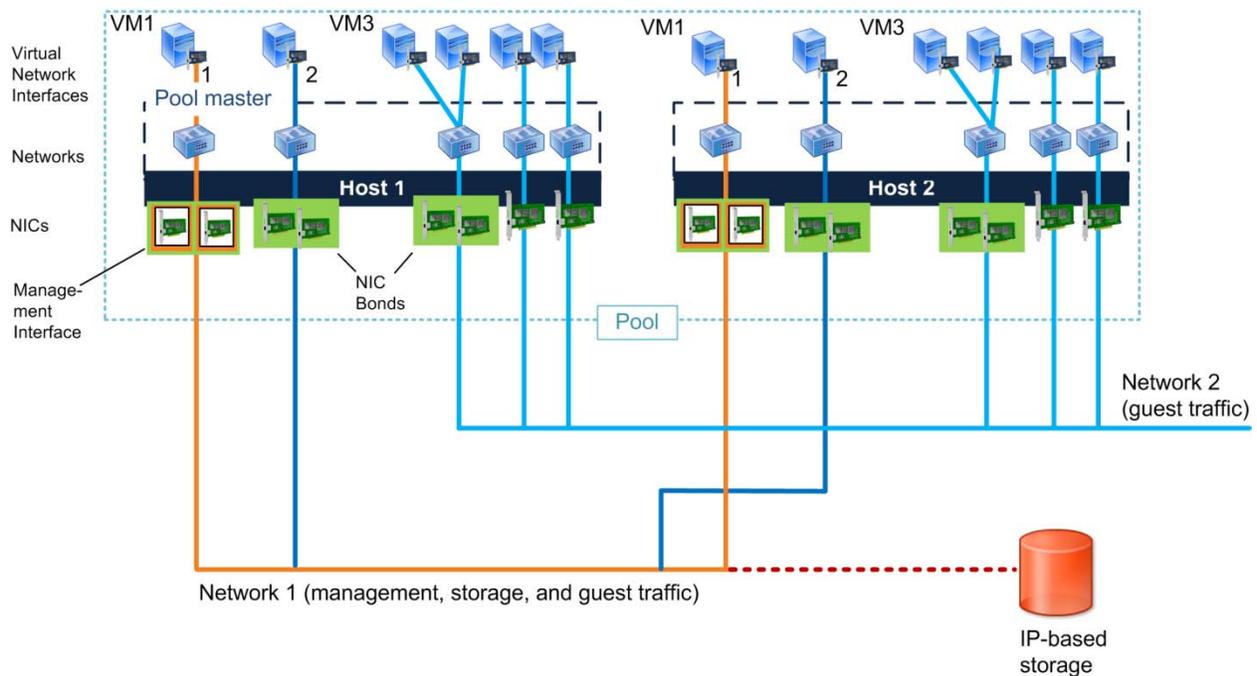
- VM traffic is also sent on this management network. This is the default configuration and requires no changes. To revert to this configuration, create a virtual network interface on the VM and specify the VM network that is sharing the management network.

This configuration lets (1) XenServer use the NIC configured as the management interface to communicate with other hosts and (2) VMs transparently forward guest traffic onto that network and back.

However, this configuration has security implications. Workstations hosting XenCenter and XenServer hosts using this management network can communicate with each other because they are on the same network. This makes the management network, which ultimately manages the hardware layer and controls the hypervisors themselves, vulnerable to any attacks originating from the VMs. For example, if the VMs host Web servers, any successful attacks originating from outside the organization can potentially penetrate your entire virtual infrastructure – or all infrastructure on the targeted pool.

In contrast, scenario 1 on page 28 separates the VM traffic from the management network, which confines any successful external attacks to the guest network.

The following illustration shows some VMs sending their VM traffic over the management network.



This logical illustration shows how the administrator configured the virtual network interfaces on VM 1 and VM 3 to send their traffic across the management network.

Note: Virtual network interfaces appear differently in Linux and Windows VMs:

- In a Windows VM, the initial Windows installation has an emulated network device that uses a built-in driver.
- In a Linux VM, the NIC appears as a standard Linux network device and uses the high-speed Xen paravirtualized network driver.

After you install the XenServer Tools (for Windows guests), Windows also uses high-speed paravirtualized network drivers.

Scenario 3: Isolating VM Traffic on a Private Network

You might have specific types of workloads that require isolation. For example, in environments with technically savvy workers, you might not want servers with confidential employee data on the same network as regular VM traffic. XenServer lets you segregate traffic by creating two types of private networks: single-server private networks and cross-server private networks.

Private networks do not have an uplink or a physical NIC. Private networks connect VMs on the same XenServer host or the same resource pool. In a private network, VMs can only communicate with VMs on the same switch on the same host. In the case of cross-server private networks, VMs can only communicate with VMs on the same vSwitch.

Essentially, a private network functions like an isolated local area network that is local to either a host or a group of hosts (pool). This results in higher speed networks since responses between VMs are based on the storage speed and not limited by the network bandwidth or bottlenecks.

Due to the speed, lab machines and test environments are a good use case for private networks. Creating private networks might also be desirable for these reasons:

- **Security.** Single-server and cross-server private networks can let you isolate VMs from other network traffic (almost like creating a virtual “stove pipe”). Private networks and cross-server private networks are completely isolated from regular network traffic. VMs outside of the private network cannot sniff or inject traffic into the network, even if both sets of VMs are on the same physical server and the virtual network interfaces on both sets of VMs transmit traffic across virtual network interfaces connected to a network on the same underlying NIC.
- **Faster traffic for connections between VMs on the same host.** Because VMs do not need to interact with regular network and switches, they can transmit traffic faster to each other.

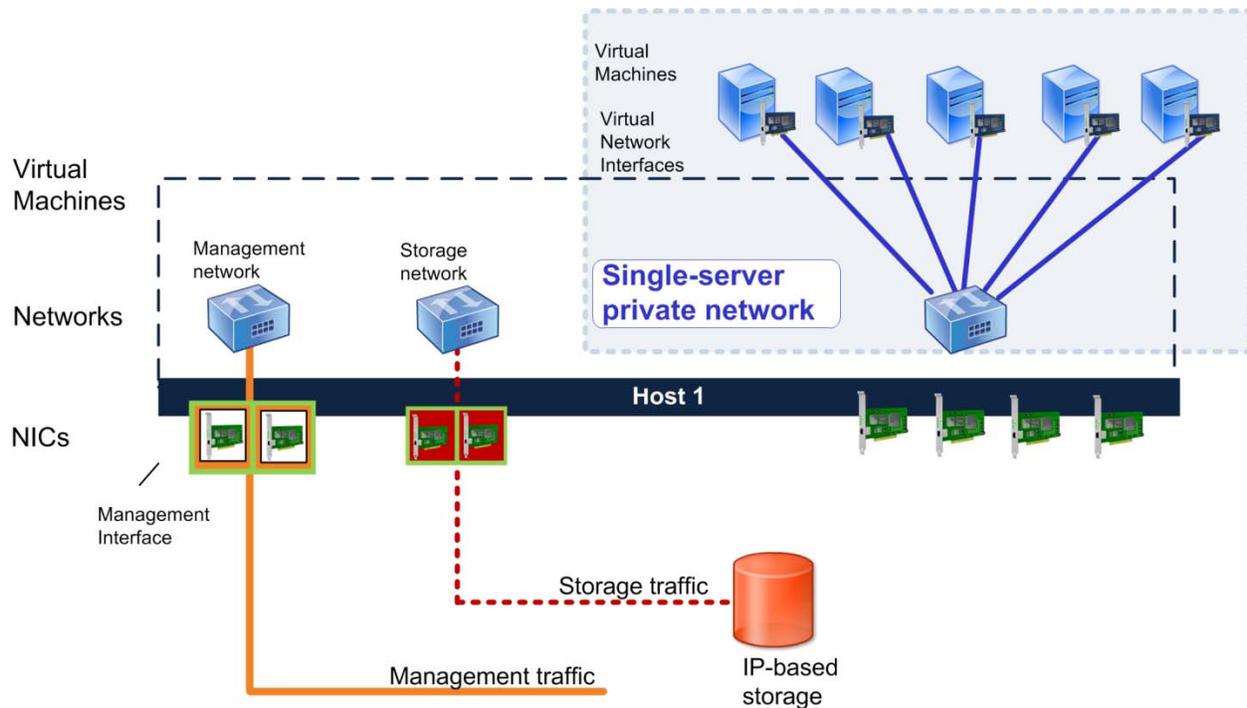
Private networks provide connectivity only between VMs on a given XenServer host and do not have a connection to the outside world. Networks with a NIC (PIF) association are considered external: they provide a bridge between virtual network interfaces and the NIC connected to the network, enabling connectivity to resources available through the NIC.

Note: In previous XenServer releases, single-server private networks were known as internal networks.

Scenario A: Isolating VM Traffic on One Host

If you have some VMs on one host that you do not want on your organization’s network, you can create a *single-server private network*. This is an internal network that has no association with a physical network interface. It only connects the virtual machines on the host and has no connection to the outside world.

The illustration that follows shows a private network configured on one host.



This illustration shows how the virtual network interfaces on the VMs are on the single-server private network. This network does not have any connect to any NICs since all traffic is sent inside the XenServer host.

To create a single-server private network that is isolated from the external network, you

1. Create a single-server private network in XenCenter.
In XenCenter, select the host in the Resource pane. Click the **Network** tab. Click **Add Network** and then select **Single-Server Private Network**.

Unlike when you create external networks, XenCenter does not prompt you to specify a NIC when you create private networks. This is because private networks do not require a NIC for connectivity.

2. Create a virtual network interface on each VM that specifies the new private network.

If you want to isolate the VMs' traffic completely, if necessary, remove any virtual network interfaces on the VMs that are on an external network.

Note: To create cross-server private networks, see CTX130423 - [Citrix XenServer 6.0 vSwitch Controller User Guide](#).

Scenario B: Isolating VM Traffic on Cross-server Private Networks

Cross-server private networks are similar to single-server private network except cross-server private networks let VMs on different hosts communicate with each other. Cross-server private networks combine the isolation properties of a single-server private network with the additional ability to span



hosts across a resource pool. This combination allows the use of VM agility features, such as XenMotion (live migration) and Workload Balancing (WLB), for VMs connected to those networks.

Cross-server private networks are completely isolated. VMs that are not connected to this type of private network cannot sniff or inject traffic into the network, even when the VMs share a host that has virtual network interfaces connected to two different networks that use the same NIC.

While VLANs provide similar functionality, cross-server private networks provide isolation without requiring physical-switch configuration.

Cross-server private networks provide the following benefits:

- The isolation properties of single-server private networks
- The ability to span a resource pool, enabling VMs connected to a private network to live on multiple hosts within the same pool

Because cross-server private networks require a NIC with an IP address, to configure these networks you must create a secondary interface. Cross-server private networks can use any secondary interface as the underlying network transport. However, if you choose to put cross-server-private network traffic on a secondary interface, this secondary interface must be on a separate subnet.

To create a cross-server private network, the following requirements must be met:

- All hosts in the pool must be using XenServer 5.6 Feature Pack 1 or greater
- All hosts in the pool must be using the vSwitch for networking
- The pool must have a vSwitch Controller configured
- The cross-server private network must be created on IP-enabled NICs (that is, the management interface or a secondary interface)

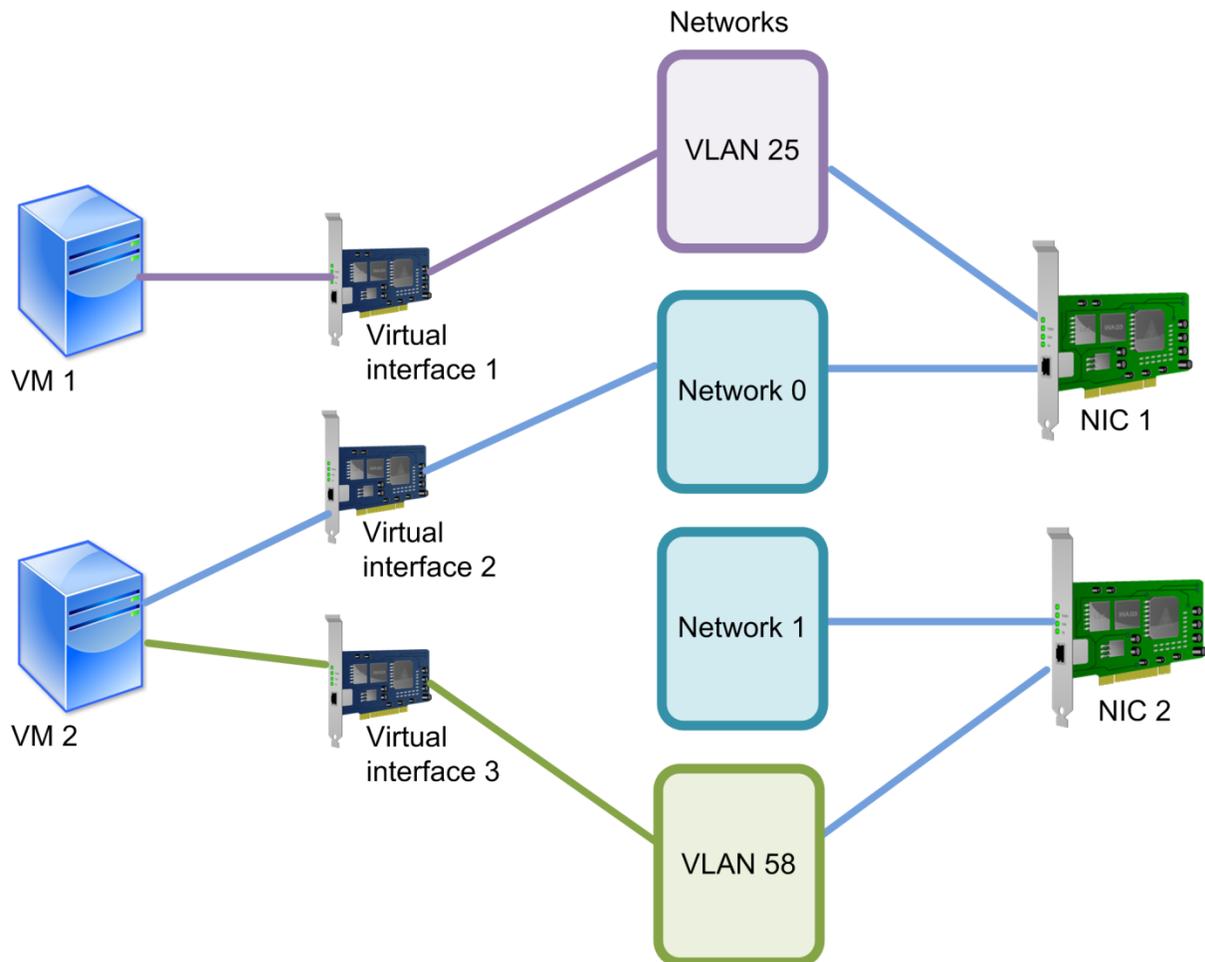
To create a cross-server private network using XenCenter

1. In XenCenter, select the pool where you want to create the network in the Resources pane and click **Add Network**.
2. On the first page of the **New Network** wizard, select **Cross-Server Private Network** and click **Next**.
3. Enter a name and description for the new network, and click **Next**.
4. Do one or more of the following:
 - To automatically add the new network to any new virtual machines created using the New VM wizard, select the check box.
 - To use jumbo frames, set the Maximum Transmission Unit (MTU) to a value between 1500 to 9216.
5. Click **Finish** to create the new network.

Scenario 4: Connecting VMs to Multiple Linked VLANs

Many organizations today configure VLANs to logically separate their physical networks for either performance or security reasons. If your organization has VLANs, you might want to connect your VMs to one or more VLANs on your network.

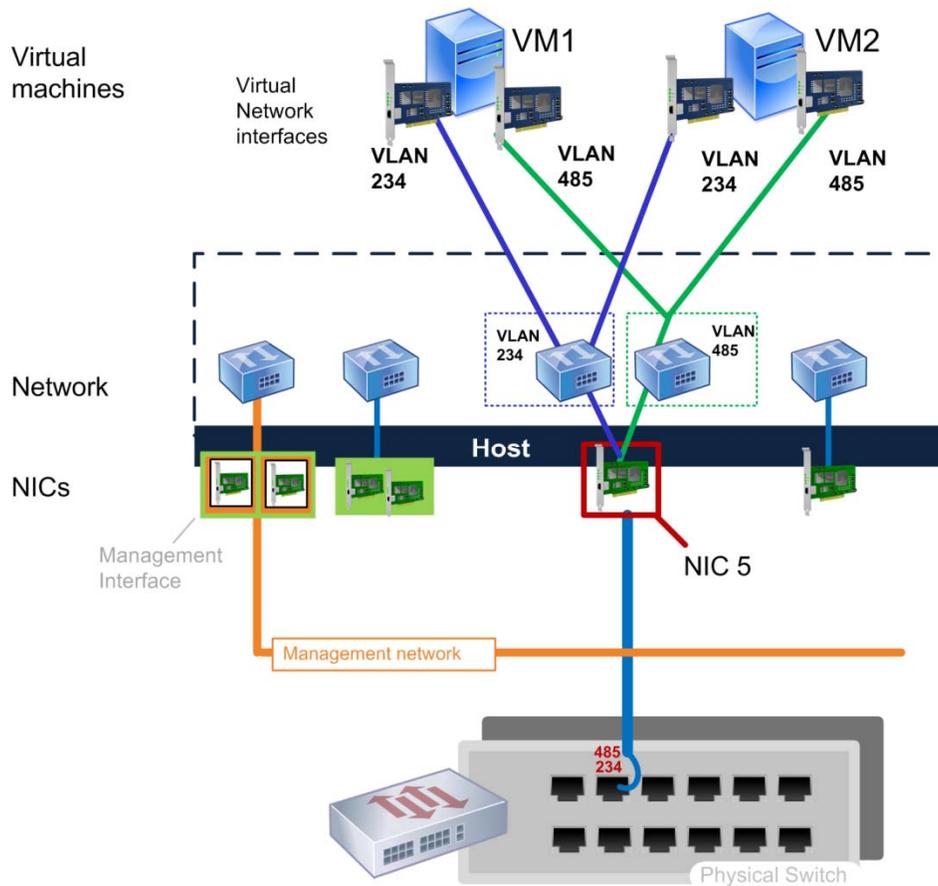
To connect a VM to a VLAN, you must create a network for the VLAN and then connect the VM to that network. To perform this configuration, you create a separate external network for each VLAN and then create a virtual network interface on the VM for each of these networks.



This illustration shows how VMs require a separate virtual network interface for each network to which you want to connect them, including VLANs. In this example, VM 2 connects to Network 0 through Virtual Network Interface 2 and to VLAN 58 through Virtual network interface 3. As shown by VM1 and NIC1, multiple networks can connect out through one NIC.

While trunk lines from the physical switch can contain multiple 802.1q VLANs, XenServer does not let you combine multiple VLANs in one XenServer network. This means that to let a VM connect to multiple VLANs you must either (a) create a separate network in XenServer for each VLAN or (b) create a XenServer network for a VLAN that can access all of the desired VLANs.

In the illustration that follows, the VMs connect to a VLAN through a trunked switch port.



This illustration shows how VMs on the host connect to an external network that the administrator configured to connect to VLAN 485 and VLAN 234. To achieve this, the administrator created an external network that uses NIC 5 to connect to a trunked switch port that includes VLAN 485 and a second external network that also uses NIC 5 to connect to VLAN 234. The administrator ran a cable from the VLAN trunk port to NIC 5.

Connecting a VM to a VLAN requires that you:

1. Create a physical connection between the corresponding NIC on each host and the VLAN trunk port for that VLAN on the switch.

For example, if you connect NIC 7 on the XenServer pool master to a VLAN trunk port on the switch with access to VLAN 485, you must run a cable from NIC 7 on all other hosts in the pool to a similarly configured VLAN trunk port on the same switch, which can access VLAN 485.

2. Enable XenServer to connect to a specific VLAN on the switch by creating an external network specifying that VLAN tag.

This means creating an external network on the XenServer pool master and specifying the VLAN tag when you create the network.

In XenCenter, select the pool (<*your-pool-name*>) in the **Resource** pane, click the **Network** tab, and click the **Add Network** button. In the New Network wizard, select **External Network**. On the **Location** page, specify the NIC you physically connected to the switch and enter the VLAN tag for the VLAN in the **VLAN** box.

In the XenServer CLI, you can use the **pool-vlan-create** xe command to create the VLAN on all hosts in a resource pool. For more information, see the *XenServer Administrator's Guide*.

After you create the network for the VLAN on the pool master, XenServer configures the NICs on all the other hosts so that the corresponding NIC on each host

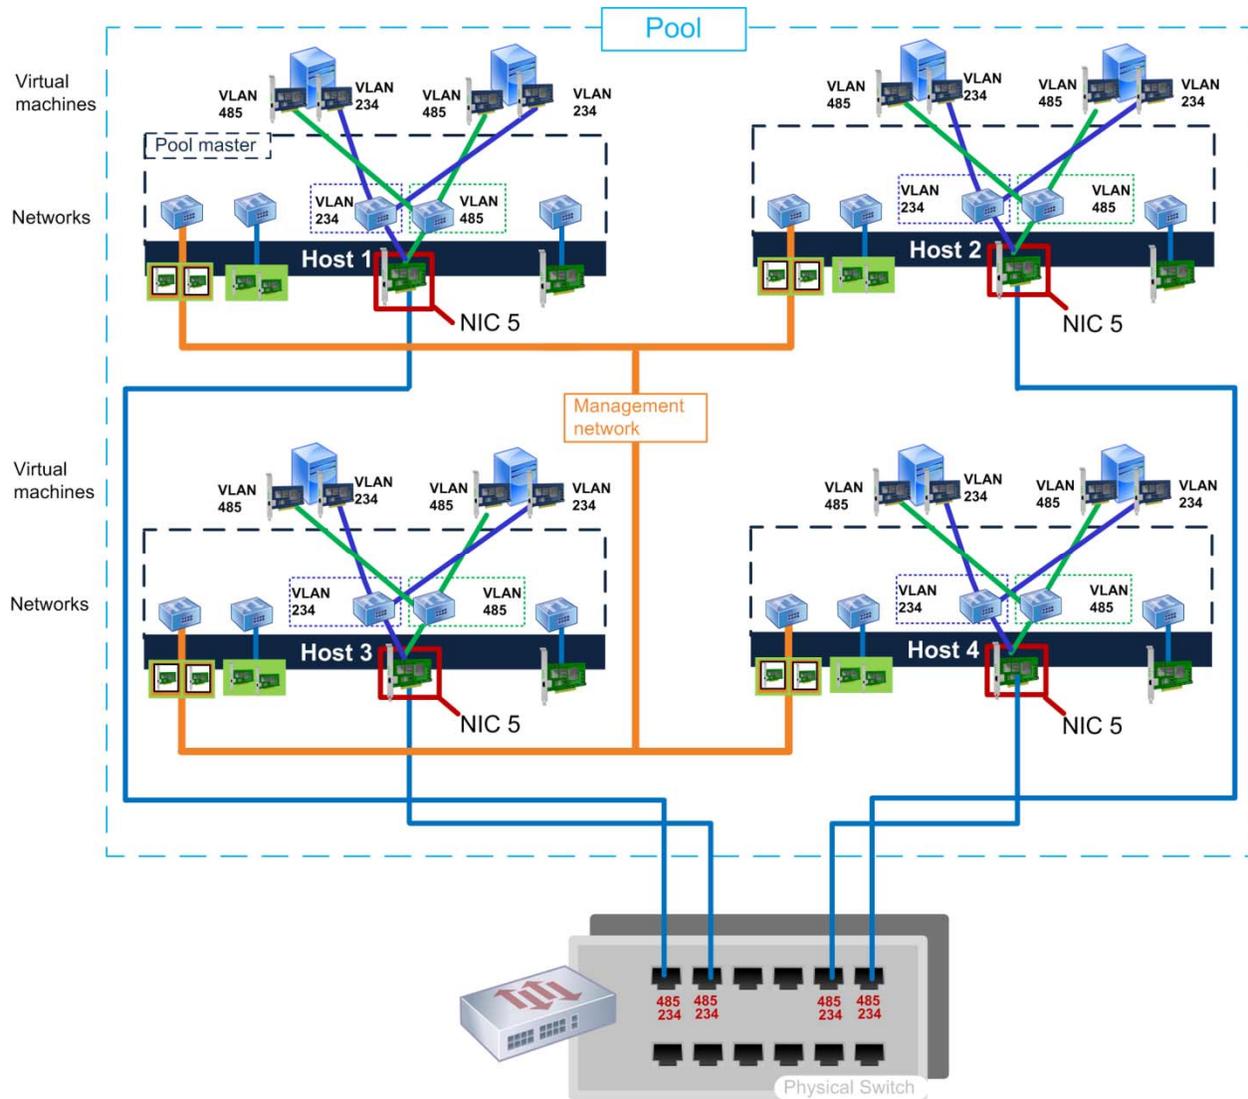
Note: The numbers of VLAN tags must be between 0 to 4094.

3. Connecting the appropriate VMs to the VLAN by configuring a virtual network interface that points to that network on each VM you want to be able to connect to the VLAN.

In XenCenter, this is done by selecting the VM in the Resource pane, clicking the **Network** tab, and clicking **Add Interface** and then specifying the VLAN network when you create the interface.

Again, because networking is a pool-level feature, if you connect one host to a VLAN, you must connect all hosts in the pool to the VLAN. This means that you must physically connect the corresponding NIC on each host to the VLAN port on the switch.

In the illustration that follows the VMs on multiple hosts in a pool connect to a VLAN through a trunked switch port.



This illustration shows how, because XenServer automatically synchronizes the network settings in pools so that they match, NIC 7 on all hosts in the pool will be configured with the same network and VLAN settings as NIC 7 on the pool master. However, for the VMs on the member servers to be able to connect to the VLAN, the administrator must also physically connect NIC 7 on each host to a trunk port on the switch that can access VLAN 485.

Before configuring a VLAN, ensure the switch on your VLAN network is configured as follows:

- The port on the switch connected to each XenServer host must be configured as trunk port.
- The port on the switch must be configured for 802.1q encapsulation.
- Port security cannot be set on the trunk port.
- The port designated as trunk should be assigned a native VLAN; use 1 as default.

XenServer lets you create multiple networks and VLAN networks on the same NIC. XenServer does not limit the number of VLANs you can connect to VMs. Instead, the limit comes from the



802.1q standard is 4096. You add an external network for each VLAN to the host and then connect the VMs to the VLANs by specifying that network in the VM's virtual network interface.

Note: If a Native VLAN is used on the switch trunk port, then you cannot assign that VLAN number to a VM on the XenServer.

For an example of a tested working model of a VLAN configuration, see CTX123489—[XenServer VLAN Networking](#). For more information about configuring VLANs on your switch and 802.1q support, see the documentation for your switches.

Tip: To verify that you have configured the XenServer host to communicate across the correct network, you can use the packet sniffing software included with your NICs to capture and display the VLAN tags that are transmitted across the switch to the XenServer.

Note: Although the 802.1q standard limits the number of VLANs XenServer supports, XenServer does not limit the number of XenServer networks you can configure for a NIC.

Understanding the Impact of Numerous Networks in a Pool

Having numerous connections to VLANs (for example, 100s) configured on a host creates an additional load on the Control Domain, which frequently results in reduced network performance as described.

Having numerous VLANs can also impact your host, pool, and VMs performance in the following ways:

- VM performance may degrade.
- VM network service may degrade. However, this can be due to many factors.
- Numerous VLANs can slow down certain host (XenAPI) operations, such as adding and removing networks.

In addition, various management and administration functions can become slower when there are numerous networks on pools. For example, actions like the following may take longer: joining a new host to a pool, rebooting a host; rendering charts in the Performance tab in XenCenter.

Creating VLANs on Bonded Networks

XenServer supports connecting to VLANs from bonded NICs. To do so, do the following:

1. Bond the two or more NICs together. After you have done so, the NIC bond appears as a bonded network in XenCenter.
2. In XenCenter, for example, create an **External Network** specifying the following:
 - a) The VLAN's tag
 - b) The NIC bond as the NIC



You might want to name this external network the same name as the VLAN (for example, VLAN 25).

3. When you create the virtual network interface for the VM, specify the external network with the VLAN tag as the network.

Creating VLANs on the Management Interface

You can have a single VLAN on the management interface, and this VLAN can be on an access port. If you want to use a trunk, you must either define a default VLAN on that trunk and the management interface can use that or you make the port a full access port.

XenServer does not support having a VLAN trunk port on the management interface.

Chapter 4: Specifying Networking Requirements

This chapter provides information about the following:

- XenServer networking requirements and support
- A suggested process for defining your communication and hardware requirements

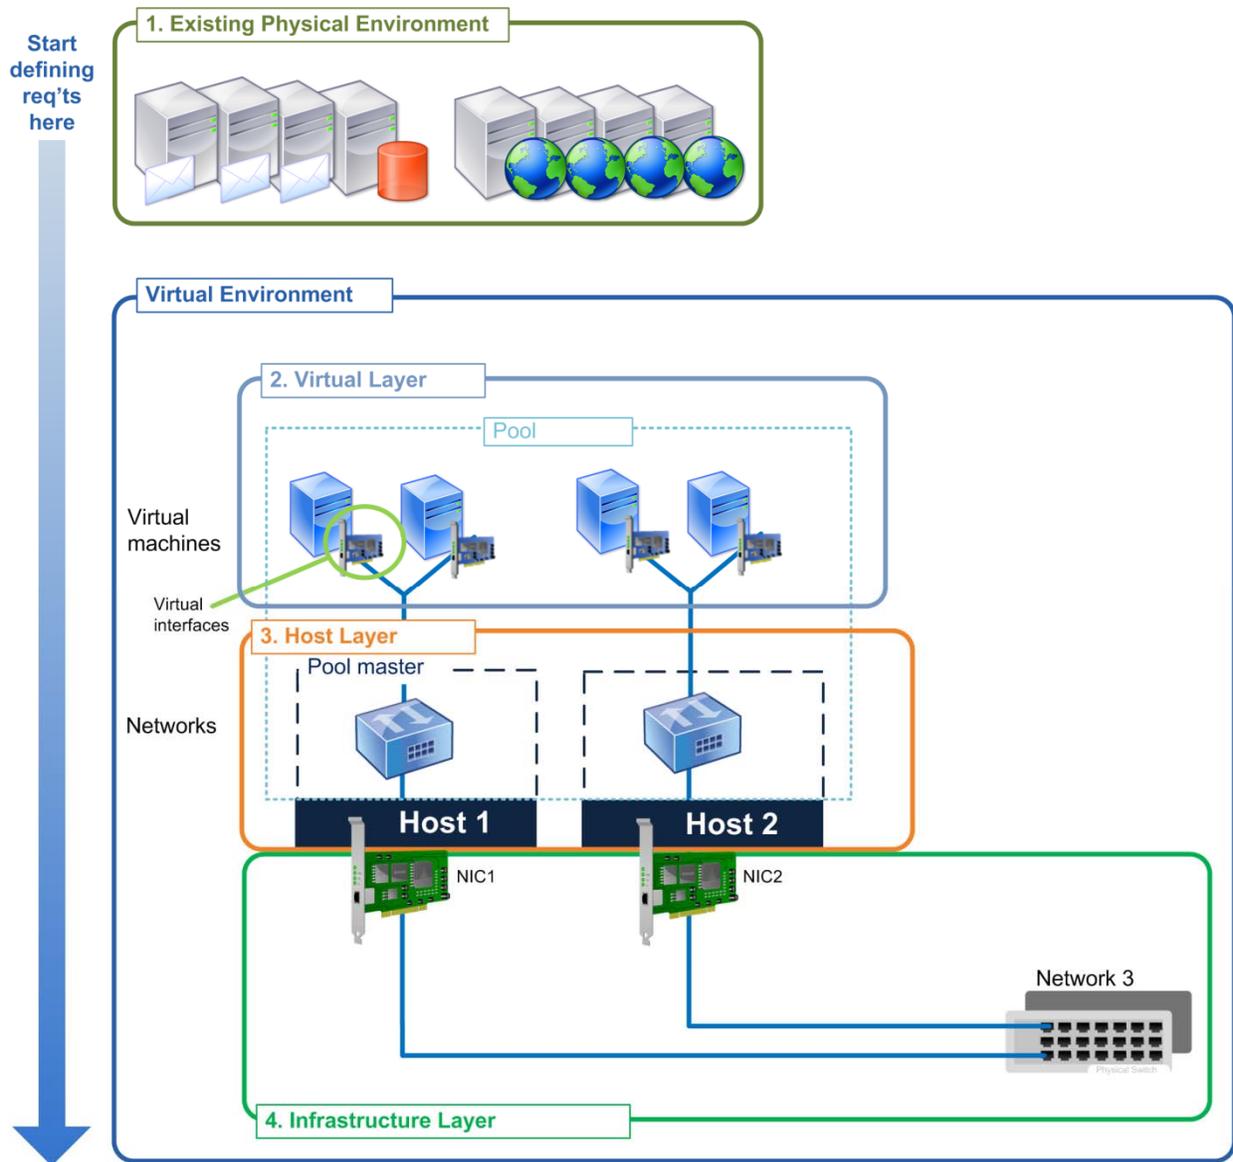
Overview

This chapter helps you map your existing networking requirements onto XenServer features and includes principles to note while defining your XenServer networking requirements. It also provides information about XenServer networking requirements and supported configurations.

This chapter includes tables that map common requirements and configuration scenarios onto XenServer features and indicate what topics to read for more information. The topics covered in the tables include guidance about defining basic requirements at the VM level and at the host level.

The tables are structured so that you begin your evaluation by considering your existing workloads (for example, in your physical environment) since this will indicate your VM's connectivity requirements. After determining your VM's communication requirements, define how you want to group VMs together and the host's networking hardware.

The following illustration shows a suggested sequence or “direction” in which to consider networking requirements: from the physical environment down to the layers of the virtual environment.



This diagram presents a possible way of defining networking requirements for a XenServer pool. It suggests that you begin by examining your physical environment and then consider your networking requirements by beginning at the VM workload level, defining networks required at the host level, and then defining your physical infrastructure requirements, such as NIC and switch configuration.

Introduction

The primary factor that drives your networking requirements is the connectivity needs of the pool's VMs and their workloads. In some cases, you may choose to group VMs in pools according to the networking requirements of their workloads. This could either be due to the:

- Workloads' networking hardware requirements since all hosts in a pool must use the same networking hardware



- Workloads' networking configuration since sometimes you might want to create pools based on common networking requirements so as to reduce the amount of networking configuration you must perform

As a result, it helps to know the workloads you want to virtualize and their networking requirements before you begin configuring networking. Likewise, you should know the approximate size of your pool (that is, the number of hosts).

Ideally, all networking configuration is performed before you put a pool into production. However, you do not need to configure all of the pools in your virtual environment when you begin configuring networking. Rather, you can configure one pool at a time. Although configuring networking before putting a pool into production is a best practice, you can add hosts and make networking changes at any time.

XenServer Networking Support and Requirements

This section provides information about the physical and logical networking configurations XenServer supports, such as the number of NICs or networks supported. It also provides information about where to find a list of supported networking hardware.

When you are defining networking requirements for a pool, it is important to note that all pooled hosts should have the same number and model of NICs, same XenServer networks, and physical cabling configuration. Because XenServer assumes all network settings in a pool match, it automatically propagates any changes you make to network settings on one host to all other hosts in the pool.

XenServer Supported Configurations

XenServer supports the following networking configurations:

- Up to 16 physical network interfaces (or up to 8 pairs of bonded network interfaces) per XenServer host
- Up to 7 virtual network interfaces per VM
- Active-active and active-passive bonding modes are supported. DMP and RDAC MPP multipath handlers are also supported
- Four different types of networks: external, cross-server private, single-server private, and VLANs
 - There is no Citrix-imposed preset limit on the number of VLANs
- SR-IOV provided the NIC used meets the support requirements in the [Citrix XenServer Hardware Compatibility List](#)



XenServer Networking Hardware Requirements

Citrix Technical Support only provides support for hardware, including NICs, on the [Citrix XenServer Hardware Compatibility List](#). While it may be possible to use different hardware, it is not recommended. If you do not see your hardware on the list, you can use the self-certification kits at the [Citrix XenServer Hardware Compatibility List](#) web page and certify it yourself by running the kit and submitting your results.

Important: Do not configure any networking settings until you have added all the hosts to the pool and you finish physically connecting each host to the appropriate switch ports. Then, proceed to configure your XenServer network settings by starting with (and configuring only) the pool master.

Defining Your Networking Requirements

Defining network requirements for your XenServer environment is a multi-stage process. Two factors significantly influence your network design and requirements:

1. Certain workloads may have specific network connectivity or performance requirements.
2. While VMs in a pool can connect to different networks, the network configurations and hardware on each host in a pool must match.

As a result, you might find it easier to define your workloads' network requirements before you determine in which servers and pools to host your VMs.

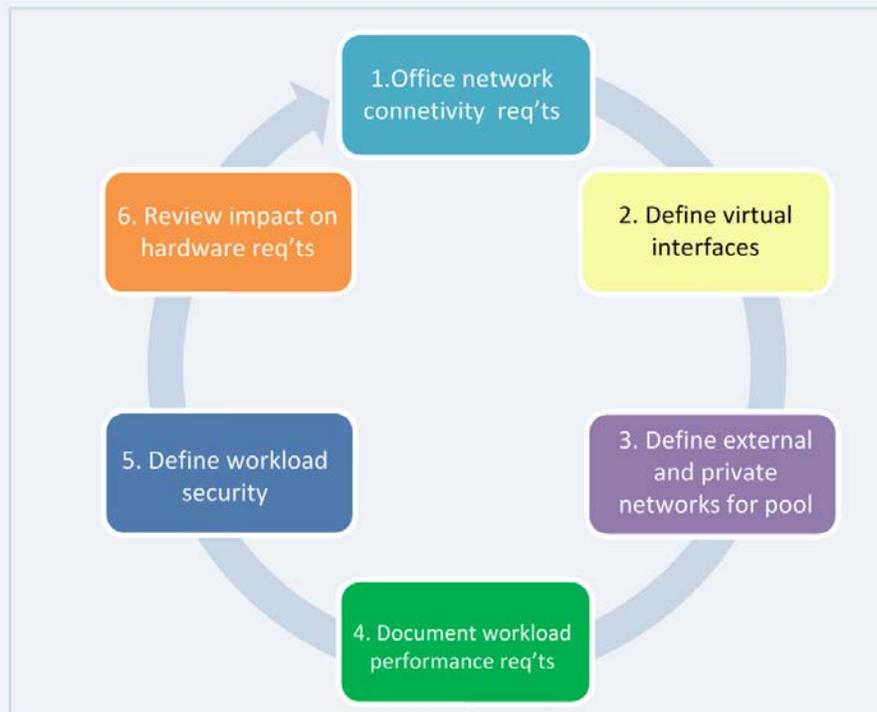
This section provides guidance about how to begin determining the communication requirements for your workloads, how to evaluate your existing networking configuration, and how to review your initial pool design against the network requirements of its workloads.

Considering Workload Communication Requirements

To determine a host's networking configuration, start evaluating requirements at the lowest level (that is, the VM/workload) and then work your way up to your organization's physical network. The networking requirements for your workloads can ultimately impact your pool design and which workloads you decide to group together.

To connect a VM to an external network, such as a VLAN, you might determine the VM's networking requirements by using a process like the one in the following illustration.

Defining Networking Requirements for a Workload



This diagram illustrates a general process you might use when defining networking requirements at the workload level.

Determine the workload’s communication requirements (for example, does a domain controller need access to a specific VLAN or a database server need access to a specific storage device?). To do so, consider the factors listed in the following table:

Factor	Action, Notes
Determine if the workload has any specific performance requirements that might change its host’s hardware design.	Some workloads might require specific NICs due to performance requirements. An option for optimizing virtualized Provisioning Services servers is SR-IOV – see “Virtualizing the Provisioning Services Server” on page 88. See the XenServer Hardware Compatibility List for supported NICs.
Consider the redundancy	For information about NIC bonding, see “Designing Network

Factor	Action, Notes
requirements for the workload. Is the workload mission critical?	<p>Redundancy” on page 56.</p> <p>For redundancy, you can also look at solutions by Citrix partners, such as Marathon, which provide multi-CPU full-compute failure fault tolerance.</p>

Evaluating Your Current Network Configuration

When designing your XenServer network configuration, consider the following aspects of your current physical networking configuration:

Factor	More information
Security and Isolation	
Consider your workload’s communication requirements and how you will connect its VM to the devices and/or network locations in your external network.	<p>VMs connect to networks using virtual network interfaces as described in “Chapter 3: Sample Networking Scenario.”</p> <p>If your VM must connect to a specific VLAN that is part of a trunk port, you must create a specific external network associated with the VLAN. For more information, see “Scenario 3: Isolating VM Traffic on a Private Network” on page 32.</p>
Does the workload need to be isolated from other traffic?	<p>Configure either a:</p> <p>VLAN subnet or a trunk port</p> <p>or a</p> <p>Cross-server private network as described on page 32.</p> <p>Additionally, you can use the Distributed Virtual Switch controller to isolate VMs by blocking specific machines or IP addresses from communicating with each other.</p> <p>Cross-Server private network is described in “Scenario 3:</p>

Factor	More information
	Isolating VM Traffic on a Private Network” on page 30.
Will you need to perform live migration for this workload?	<p>For pools with VMs that use large amounts of memory or pools that will have frequent migrations, consider using the primary management physical interface just for live migration and management tasks. In this scenario, all VM, storage, and other traffic should be placed on separate physical interfaces. See “Segregating VM Traffic from Management and Storage Traffic” on page 27.</p> <p>Live migration requires shared storage.</p>
Do you want to use four NICs in a bond?	To use four NICs for resiliency, you must use the vSwitch as your networking stack. For a description of the different network stacks, see “Deciding to Use the Distributed Virtual Switch” on page 53.
Configuration	
Do you want XenServer to automatically generate MAC addresses for the VM hosting this workload or do you want to assign the VM a static IP address?	Citrix generally recommends letting XenServer automatically generate MAC addresses. See “Understanding Virtual MAC Addressing” on page 26.

Determining Host Networking Requirements

Factor	More information
Does the workload need all of the bandwidth available or should it be deprioritized in favor of other workloads?	Set a quality of service (QoS) restriction on the virtual network interface for the VM. See “Limiting Bandwidth Consumption for High Demand Workloads” on page 72.
Do you intend to perform storage migration (Storage XenMotion) frequently? Do your virtual disks or virtual machines contain a lot	Consider configuring storage migration to occur on its own network and NIC to offset the impact of migration traffic on the management network.

Factor	More information
of data?	To do so, you must assign an IP address to the NIC you want to dedicate to storage motion. See “Creating a Dedicated Network for Storage Motion” on page 74.
Storage Requirements	
Does the storage device require an IP address on the host?	If so, you may want to configure a secondary interface so you can assign an IP address to the NIC. For more information, see “Segregating VM Traffic from Management and Storage Traffic” on page 27 and “Chapter 6: Designing Your Storage Network Configuration” on page 76.
Are you planning to use an iSCSI Host Bus Adapter (HBA)?	For a list of supported HBAs, see the Citrix XenServer Hardware Compatibility List .
Would the storage traffic benefit from configuring jumbo frame support?	<p>Currently, jumbo frames are only supported for storage networks with iSCSI HBAs and the vSwitch configured as the networking bridge. This means:</p> <ol style="list-style-type: none"> 1. If you want to configure jumbo frames for your storage traffic, you must use a storage device that can use an HBA, such as an iSCSI hardware SAN or a Fibre Channel SAN. 2. You must configure the vSwitch as the networking bridge. You can choose to configure the Distributed Virtual Switch solution or just configure the vSwitch as the bridge. See “Deciding to Use the Distributed Virtual Switch” on page 53. 3. You must configure end-to-end support for your jumbo frames, including switches and NICs that are compatible with them. For more information, see “Configuring Networks with Jumbo Frames” on page 85.
Do you need to change duplex settings on the host?	See CTX117568— How to Modify Network Speed and Duplexing .



Reviewing Initial Pool Design against Networking Requirements

After you have determined a possible pool design for a group of workloads, consider if any of the following will cause problems:

Factor	Action, Notes
Will the network hardware requirements for these workloads clash? Hardware in pools should match.	If a subset of workloads require more expensive NICs, such as ones that support jumbo frames or 10 gigabit Ethernet, do you want to purchase that networking hardware only for the servers hosting those workloads? If so, you should group these workloads in the same pool(s). See “Impact of Pools on XenServer Networking” on page 15.

Considering Addressing Requirements

This section discusses IP addressing considerations.

Factor	Action, Notes
How do you want to configure IP addressing for the primary and secondary interfaces?	<p>Unless you choose to add an IP address to a NIC, only the management interface requires an IP address: by default, all other NICs do not have an IP addresses.</p> <p>You can assign static IP addresses to each NIC in the host through XenCenter or the xsconsole.</p> <p>You specify the IP address for the management interface during XenServer Setup. You can specify that XenServer uses either static or dynamic IP addresses. However, assigning hosts static IP addresses is generally preferred. For more information, see “Networking Configuration after Installation” on page 14.</p> <p>IP-based storage requires configuring a secondary interface so you can assign an IP address to the NIC. For more information, see “Segregating VM Traffic from Management and Storage Traffic” on page 27 and “Chapter 6: Designing Your Storage Network Configuration” on page 76</p> <p>For information about setting a static IP address after Setup, see CTX116372—How to Assign a Static IP Address to a XenServer Host.</p>

Calculating the Number of Physical NICs per Host

All XenServer hosts in a pool should have the same number of NICs; however, this requirement is not strictly enforced when a XenServer host joins a pool.

Having the same physical networking configuration for XenServer hosts within a pool is important because all hosts in a pool share a common set of XenServer networks. The NICs on each host connect to pool-wide networks based on device name. For example, all XenServer hosts in a pool with an eth0 NIC will have a corresponding NIC plugged into the pool-wide Network 0 network. The same will be true for hosts with eth1 NICs and Network 1, as well as other NICs present in at least one XenServer host in the pool.

If one XenServer host has a different number of NICs than other hosts in the pool, issues can occur because not all pool networks will be valid for all pool hosts. For example, if host1 and host2 are in the same pool and host1 has four NICs while host2 only has two, only the networks connected to NICs corresponding to eth0 and eth1 are valid on host2. VMs on host1 with virtual network interfaces connected to networks corresponding to eth2 and eth3 will not be able to migrate to host2.

The number of physical NICs you want on each host (and, consequently, the pool) depends on your required resiliency and connectivity. For an example of a reason you might want to use additional NICs to separate traffic, see “Segregating VM Traffic from Management and Storage Traffic” on page 27.

Although XenServer can be run with only one NIC on the host, Citrix recommends having at least two NICs on the host: one for VM traffic and one for management traffic. Other examples of how you could use NICs include:

- A best practice is to bond the NICs for the management interface and the one for VM traffic, which means deploying at least four NICs per host. If you are using IP-based storage, you might want six NICs.
- You may want to provide additional bandwidth for VM traffic by adding additional NICs to the host.
- If you are going to have shared storage, consider dedicating a physical network for your storage traffic. In this case, consider having at least one NIC or HBA, or ideally two in a bonded or multipathed configuration, dedicated to the storage traffic.
- If you are using Provisioning Services to stream disk images to VMs, you may want to dedicate a bond dedicated for the Provisioning Services server on the XenServer host.

Consider the factors in the following table to help determine the number of NICs you want in your hosts:

Factor	Action, Notes
Do you want to optimize your IP-based storage or provide redundancy?	<ul style="list-style-type: none"> • The best practice is to configure two NICs or iSCSI HBAs for storage traffic in a bonded or multipathed configuration. • For bonding, see: <ul style="list-style-type: none"> ○ “Creating Network Resiliency through Bonds” on page 23. ○ “Considering NIC Bonding” on page 56. • For multipathing, see: <ul style="list-style-type: none"> ○ CTX121364—Multipath Implementations in XenServer. ○ CTX129309—Configuring iSCSI Multipathing Support for XenServer ○ CTX118791—Multipathing Overview for XenServer 5.0. This article provides a good, albeit somewhat dated, overview of the UI-based multipathing configuration process.

For information about adding additional NICs to XenServer, see CTX121615—[How to Add Additional Physical NICs to XenServer](#).

For an example of how to calculate the number of NICs for XenDesktop-XenServer deployments, see the 22 February 2011 Citrix blog post, “[XenServer for XenDesktop - How many network cards do I need?](#)”

Calculating Bandwidth Requirements

Estimating how much network bandwidth your virtualized environment will need is a key part of ensuring good VM performance. Because all VMs on a host share the host’s bandwidth, providing the host with enough bandwidth is critical. Factors that affect bandwidth requirements include:

Factor	Action, Notes
Number of VMs	<ul style="list-style-type: none"> • Servers hosting more VMs may require more bandwidth, depending on the type of workload. See “Designing Networks for Performance” on page 67.

Factor	Action, Notes
Type of workload and traffic	<ul style="list-style-type: none"> • Some server roles require more bandwidth. For example, servers sending a lot of traffic to storage devices or that have a lot of back up traffic. • Some operating systems have a lower impact on network performance. See “Designing Networks for Performance” on page 67.
Workload-specific bandwidth requirements	<ul style="list-style-type: none"> • In some cases, you might need to constrain VMs to ensure sufficient bandwidth. For example, you might want to configure QoS policies on the NIC to specify a rate limit. See “Limiting Bandwidth Consumption for High Demand Workloads” on page 72.
Provisioning Services	<ul style="list-style-type: none"> • A key concern in a Provisioning Services deployment, especially for large XenDesktop implementations, is the IOPS required for servers and target devices when the VMs boot. See CTX128645—<i>Design Considerations for Virtualizing Provisioning Services</i>.

For information about testing XenServer network performance, see “Resolving Network Bottlenecks” on page 69.

Chapter 5: Designing XenServer Networks

This chapter provides an overview of the decisions you need to make when designing your XenServer networking configuration and contains information about the following:

- Deciding whether or not to implement Distributed Virtual Switching
- Designing your XenServer networking configuration for redundancy
- Isolating traffic
- Designing for performance
- Hardware and NIC specific configurations

Overview

At a high level, most XenServer network design decisions stem from a combination of three major design goals: the need for redundancy, performance, or isolation. For many organizations, these goals may overlap and are not necessarily mutually exclusive.

While considering these goals, keep in mind that the physical network configurations you create for one host should match those on all other hosts in the pool. Consequently, it helps to understand similarities between networking requirements for different workloads before you group workloads together in hosts and, ultimately, pools.

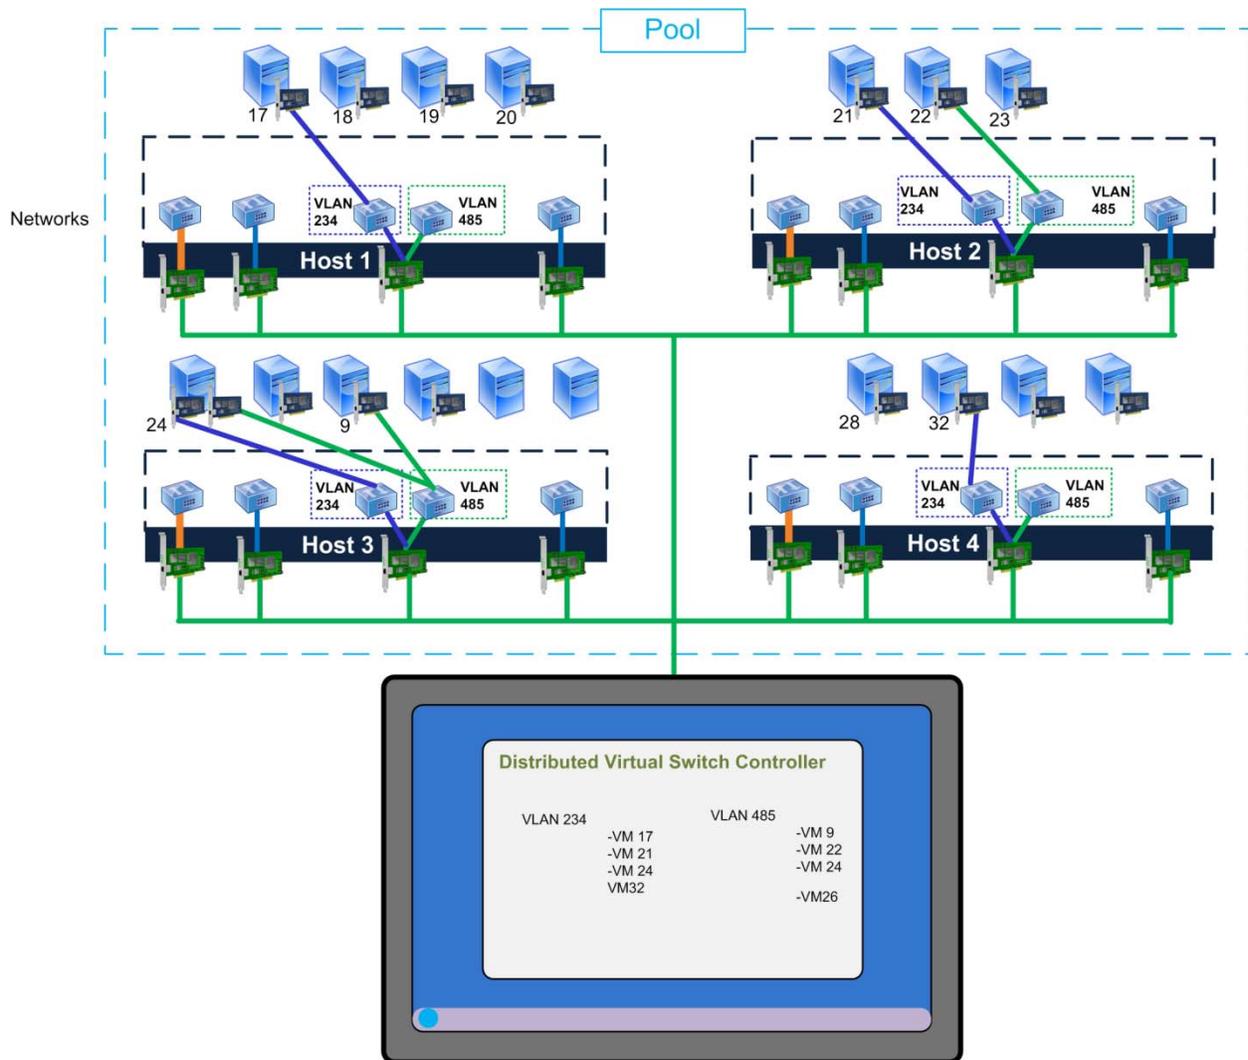
When designing your overall network configuration, you must determine your bonding and management-interface requirements; internal, external, and storage networking requirements; and VLAN configuration. You must also consider your IP addressing and DNS configuration. In addition, you should consider your NIC hardware configuration and other NIC settings, such as quality of service restrictions.

The sections that follow consider these decision points as they relate to the three primary design goals of redundancy, performance, and isolation.

Deciding to Use the Distributed Virtual Switch

As of XenServer 6.0, the new XenServer vSwitch network stack is the default networking configuration. However, you can still use the Linux bridge, which was the default network stack prior to XenServer 6.0, by running an XE command to change your network stack.

Citrix recommends that customers creating new XenServer pools use the XenServer vSwitch network stack to prevent the need for future reconfiguration. Citrix is discontinuing active development on the Linux bridge and may discontinue support for the Linux bridge in future releases.



This illustration shows how the vSwitch Controller can display which VLANs are used by which virtual machines and let you display these VLANs from one central user interface.



LACP bonding requires the vSwitch network stack.

The Distributed Virtual Switch solution provides:

- Isolation through features such as Cross-Server Private Networks.
- Quality of Service (QoS) policies.
- Jumbo frame support.
- Fine-grained security policies. You can create access control lists by using the vSwitch Controller and restrict certain types of traffic to specific VMs.
- A central management console to manage finer grained features and monitor traffic. For the switch port associated with each VM, you can see the packets traversing that switch port.
- Visibility into XenServer networking through standard tools and protocols, such as RSPAN and NetFlow.
- Simplified administration of virtualized networking environments.

How does it work?

The Distributed Virtual Switch is a solution based on the Open vSwitch, an open-source project.

The Distributed Virtual Switch comprises two networking components:

- **XenServer Open vSwitch.** The XenServer Open vSwitch, or “vSwitch,” is the actual networking component running on each XenServer host. The XenServer Open vSwitch is a virtualization-aware switch. This switch is referred to as an *Open vSwitch*.
- **Distributed vSwitch Controller.** The Distributed vSwitch Controller is console on a centralized server, which is distributed as an appliance, that manages and coordinates the behavior of each individual Open vSwitch to provide the appearance of a single distributed virtual switch. If you want to manage all the vSwitches on your hosts centrally and have them function as a single switch, you must download the Distributed vSwitch Controller appliance.

From a conceptual perspective, the vSwitch functions the same way as the existing Linux bridge. Regardless of whether or not you use the Distributed Virtual Switch or the Linux bridge, you can still use the same networking features in XenCenter and the same xe networking commands listed in the *XenServer Administrator's Guide*. In the diagrams throughout this guide, if you replace the existing “network” icons, which represent the Linux bridge, with a vSwitch, the concepts remain the same.

Because networking is a pool-level feature, if you want to use the Distributed Virtual Switch solution, you must configure all hosts in the pool to do so.



Configuring the Distributed Virtual Switch Solution

Citrix recommends setting up the networking configuration you want (vSwitch or Linux bridge) before you put your pool into production. The vSwitch is the default networking configuration, but the Distributed Virtual Switch solution, which includes the vSwitch Controller Virtual Appliance, is not configured by default.

To ensure that XenServers can always reach an active vSwitch Controller, we recommend the use of Citrix High Availability for the vSwitch Controller VM. Refer to the *XenServer Administrator's Guide* for instructions on enabling high availability. Because continuous operation of the DVS Controller is critical to the operation of networking for all virtual machines, the vSwitch Controller VM restart-priority should be set to 1 and **ha-alwaysrun** should be set to **true**.

Configuring the Distributed Virtual Switch solution requires that you download the vSwitch Controller Virtual Appliance from [My Citrix.com](http://MyCitrix.com). The vSwitch Controller Virtual Appliance lets you manage your Distributed Virtual Switch implementation. For information about using the controller, see CTX130423 - [Citrix XenServer 6.0 vSwitch Controller User Guide](#).

Note: The Distributed Virtual Switch Controller is available in Citrix XenServer Advanced Edition or higher.

Determining Which Networking Backend is Configured

If you are unsure what networking backend is configured on your hosts, you can find out by running an `xe` command.

To find out which networking backend is configured

1. Run the following command on any host in the pool:

```
xe host-list params=software-version
```

2. In the output, check the results for **network_backend**. The networking stack configured appears beside **network_backend**. For example:

```
network_backend: openvswitch
```

Reverting to the Linux Networking Configuration

If you want to revert to the legacy Linux networking bridge, run the following command on each host in the pool:

```
xe-switch-network-backend bridge
```

Reboot the host after running this command.

Warning: The Linux network stack is not open-flow enabled, does not support Cross Server Private Networks, and cannot be managed by the XenServer vSwitch Controller.



Important: Before switching networking configurations, shut down all the VMs on the host first.

Considerations:

- If you want to change the network stack, you must run the **xe-switch-network-backend** command on each host in the pool separately. The **xe-switch-network-backend** command is not a pool-wide command. This command can also be used to revert to the default vSwitch networking bridge by using this syntax: **xe-switch-network-backend openvswitch**.
- All hosts in the pool must use the same networking stack. Do not configure some hosts in the pool to use the Linux bridge and others to use the vSwitch bridge.
- When you are changing your hosts to use a different network stack, you do not need to put the hosts into Maintenance mode. You just need to run the **xe-switch-network-backend** command on each host and reboot the hosts.

Note: After changing the networking configuration, check to make sure that your bonds are still enabled.

Designing Network Redundancy

As pressure on organizations to ensure network connectivity and availability for corporate resources increases, it is becoming important to improve network resiliency by ensuring redundancy at all network infrastructure levels. Citrix recommends that you consider the following network failure points when considering redundancy:

1. Network card.
2. Network cable (for example, if it is disconnected or damaged).
3. Switch (for example, if the power supply fails). You also might need to take switches offline for planned outages, such as firmware upgrades.

XenServer provides support for NIC bonding to provide network redundancy.

Considering NIC Bonding

XenServer provides three different types of bonds, all of which can be configured using either XenCenter or the CLI:

- Active/Active mode, with VM traffic balanced between the bonded NICs.



- Active/Passive mode, where only one NIC actively carries traffic.
- LACP Link Aggregation, in which active and stand-by NICs are negotiated between the switch and the server.

Note: Bonding is set up with an Up Delay of 31000ms and a Down Delay of 200ms. The seemingly long Up Delay is deliberate because of the time some switches take to actually enable the port. Without a delay, when a link comes back after failing, the bond could rebalance traffic onto it before the switch is ready to pass traffic. If you want to move both connections to a different switch, move one, then wait 31 seconds for it to be used again before moving the other. For information about changing the delay, see the *XenServer Administrator's Guide*.

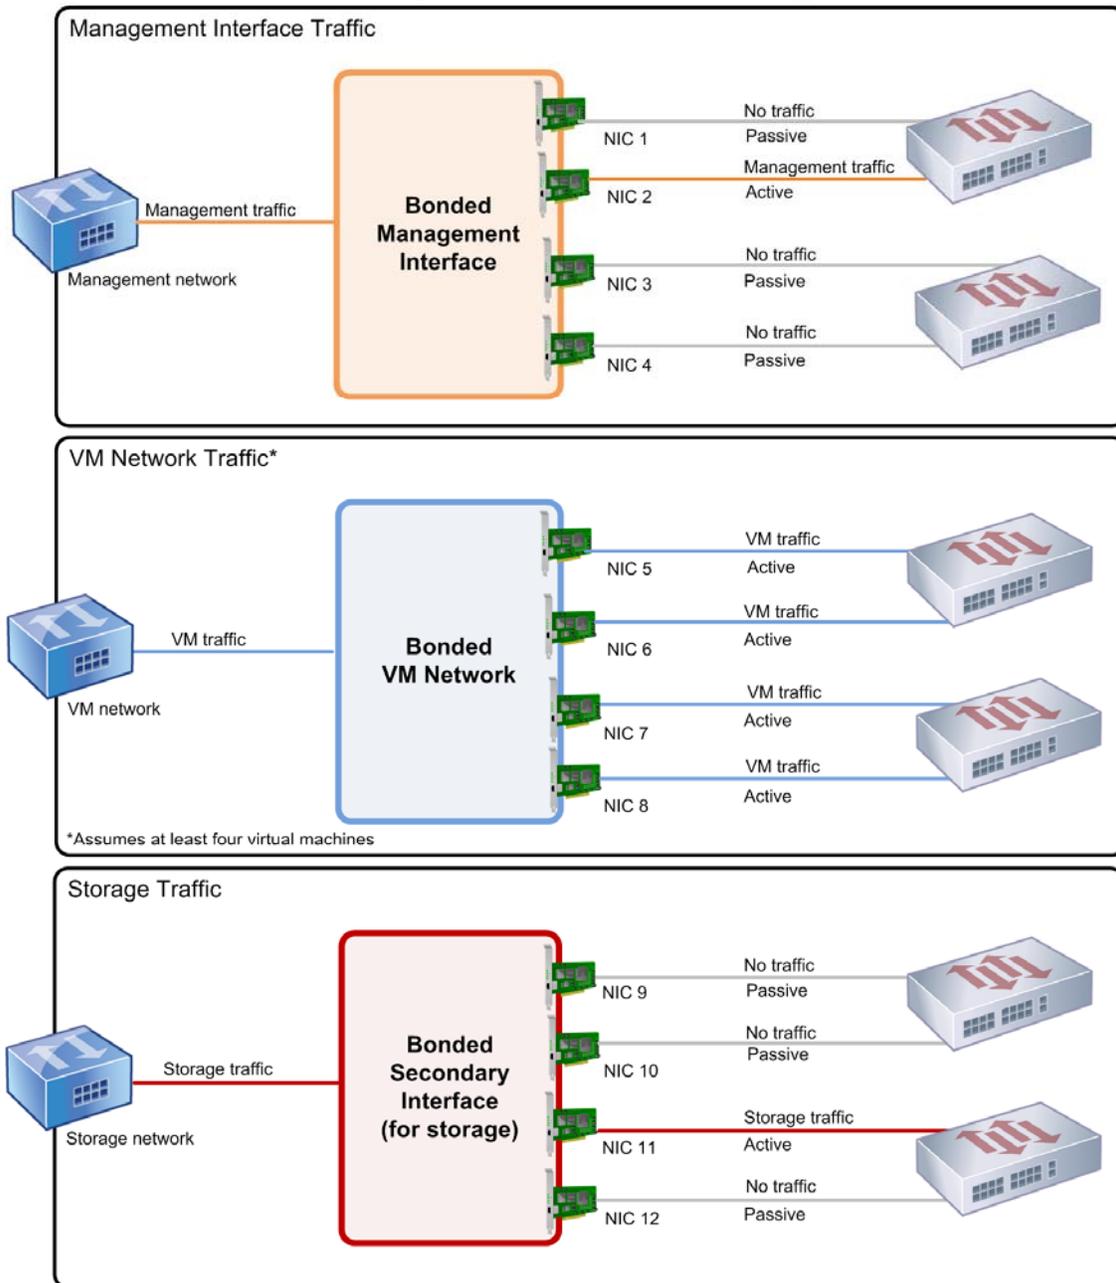
Active-Active Bonding

Active-active is an active/active configuration for guest traffic: both NICs can route VM traffic simultaneously.

When bonds are used for management traffic, only one NIC in the bond can route traffic: the other NIC remains unused and provides fail-over support. Active-active mode is the default bonding mode when either the Linux bridge or vSwitch network stack is enabled.

When active-active bonding is used with the Linux bridge, you can only bond two NICs. When using the vSwitch as the network stack, you can bond two, three, or four NICs in active-active mode. However, in active-active mode, bonding three or four NICs is generally only beneficial for VM traffic, as shown in the illustration that follows.

Active-active bonds (vSwitch network stack)



This illustration shows how bonding four NICs may only benefit guest traffic. In the top picture of a management network, NIC 2 is active but NICs 1, 3, and 4 are passive. For the VM traffic, all four NICs in the bond are active; however, this assumes a minimum of four VMs. For the storage traffic, only NIC 11 is active.

XenServer can only send traffic over two or more NICs when there is more than one MAC address associated with the bond. XenServer can use the virtual MAC addresses in the virtual network interface to send traffic across multiple links. Specifically:



- **VM traffic.** Provided you enable bonding on NICs carrying only VM (guest) traffic, all links are active and NIC bonding can balance spread VM traffic across NICs. An individual virtual network interface's traffic is never split between NICs.
- **Management or storage traffic.** Only one of the links (NICs) in the bond is active and the other NICs remain unused unless traffic fails over to them. Configuring a management interface or secondary interface on a bonded network provides resilience.

Note: Throughout this guide, the term *managed interface* is used to refer to management interfaces and secondary interfaces collectively

- **Mixed traffic.** If the bonded NIC carries a mixture of IP-based storage traffic and guest traffic, only the guest and control domain traffic are load balanced. The control domain is essentially a virtual machine so it uses a NIC like the other guests. XenServer balances the control domain's traffic the same way as it balances VM traffic.

Traffic Balancing

XenServer balances the traffic between NICs by using the source MAC address of the packet. Because, for management traffic, only one source MAC address is present, active-active mode can only use one NIC and traffic is not balanced. Traffic balancing is based on two factors:

- The virtual machine and its associated virtual network interface sending or receiving the traffic
- The quantity of data (in kilobytes) being sent

XenServer evaluates the quantity of data (in kilobytes) each NIC is sending and receiving. If the quantity of data sent across one NIC exceeds the quantity of data sent across the other NIC, XenServer rebalances which virtual network interfaces use which NICs. The virtual network interface's entire load is transferred; one virtual network interface's load is never split between two NICs.

While active-active NIC bonding can provide load balancing for traffic from multiple VMs, it cannot provide a single VM with the throughput of two NICs. Any given virtual network interface only uses one of the links in a bond at a time. As XenServer periodically rebalances traffic, virtual network interfaces are not permanently assigned to a specific NIC in the bond.

Active-active mode is sometimes referred to as Source Load Balancing (SLB) bonding since XenServer uses SLB to share load across bonded network interfaces. SLB is derived from the open-source Adaptive Load Balancing (ALB) mode and reuses the ALB capability to dynamically re-balance load across NICs.

When rebalancing, the number of bytes going over each slave (interface) is tracked over a given period. When a packet to be sent contains a new source MAC address, it is assigned to the slave interface with the lowest utilization. Traffic is re-balanced at regular intervals.

Each MAC address has a corresponding load and XenServer can shift entire loads between NICs depending on the quantity of data a VM sends and receives. For active-active traffic, all the traffic from one VM can be sent on only one NIC.

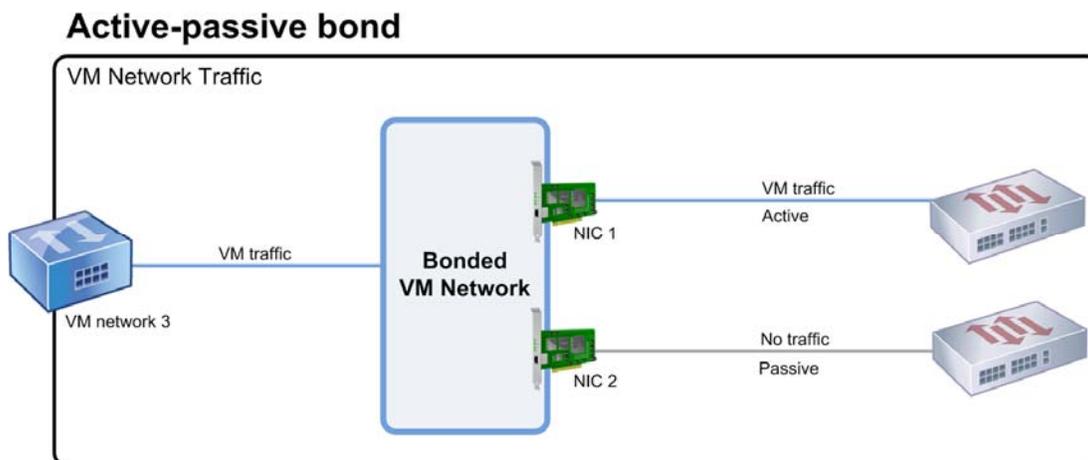
Note: Active-active bonding does not require switch support for EtherChannel or 802.3ad (LACP).

Active-Passive Bonding

An active-passive bond routes traffic over only one of the NICs, so traffic fails over to the other NIC in the bond if the active NIC loses network connectivity. Active-passive bonds route traffic over the active NIC: only if the active NIC fails is traffic shifted to the passive NIC.

Active-passive bonding is available in the Linux bridge and the vSwitch network stack. When used with the Linux bridge, you can bond two NICs together. When used with the vSwitch, you can only bond two, three, or four NICs together. However, regardless of the traffic type, when you bond NICs together in active-passive mode, only one link is active and there is no load balancing between links.

The illustration that follows shows two bonded NICs configured in active-passive mode.



This illustration shows two NICs bonded in active-passive mode. NIC 1 is active. The bond includes a NIC for failover that is connected to a second switch. This NIC will be used only if NIC 1 fails.

Since active-active mode is the default bonding configuration in XenServer, if you are configuring bonds using the CLI, you must specify a parameter for the active-passive mode or the bond is created as active-active. However, you do not need to configure active-passive mode just because a network is carrying management traffic or storage traffic.

Active-passive can be a good choice for resiliency since it offers several benefits. With active-passive bonds, traffic does not move around between NICs. Likewise, active-passive bonding lets you configure two switches for redundancy but does not require stacking. (If the management switch dies, stacked switches can be a single point of failure.)



Active-passive mode does not require switch support for EtherChannel or 802.3ad (LACP).

Consider configuring active-passive mode in situations when you do not need load balancing or when you only intend to send traffic on one NIC.

Important: After you have created virtual network interfaces or your pool is in production, be extremely careful about making changes to bonds or creating new bonds.

LACP Link Aggregation Control Protocol Bonding

LACP Link Aggregation Control Protocol is a type of bonding that bundles a group of ports together and treats it like a single logical channel. LACP bonding provides failover and can increase the total amount of bandwidth available.

Unlike other bonding modes, LACP bonding requires configuring both sides of the links: creating a bond on the host and, on the switch, creating a Link Aggregation Group (LAG) for each bond. To use LACP bonding, you must configure the vSwitch as the network stack. Also, your switches must support the IEEE 802.3ad standard.

The following table compares active-active SLB bonding and LACP bonding:

	Benefits	Considerations
Active-active SLB Bonding	<ul style="list-style-type: none">• Can be used with any switch on the XenServer Hardware Compatibility List.• Does not require switches that support stacking.• Supports four NICs.	<ul style="list-style-type: none">• Optimal load balancing requires at least one NIC per VIRTUAL NETWORK INTERFACE.• Storage or management traffic cannot be split on multiple NICs.• Load balancing occurs only if multiple MAC addresses are present.
LACP bonding	<ul style="list-style-type: none">• All links can be active regardless of traffic type.• Traffic balancing does not depend on source MAC addresses, so all traffic types can be balanced.	<ul style="list-style-type: none">• Switches must support the IEEE 802.3ad standard.• Requires switch-side configuration.• Supported only for the vSwitch.• Requires a single switch or stacked switch.

Traffic Balancing

XenServer supports two LACP bonding hashing types —the term hashing refers to the way in the NICs and the switch distribute the traffic— (1) load balancing based on IP and port of source and destination addresses and (2) load balancing based on source MAC address.

Depending on the hashing type and traffic pattern, LACP bonding can potentially distribute traffic more evenly than active-active NIC bonding.

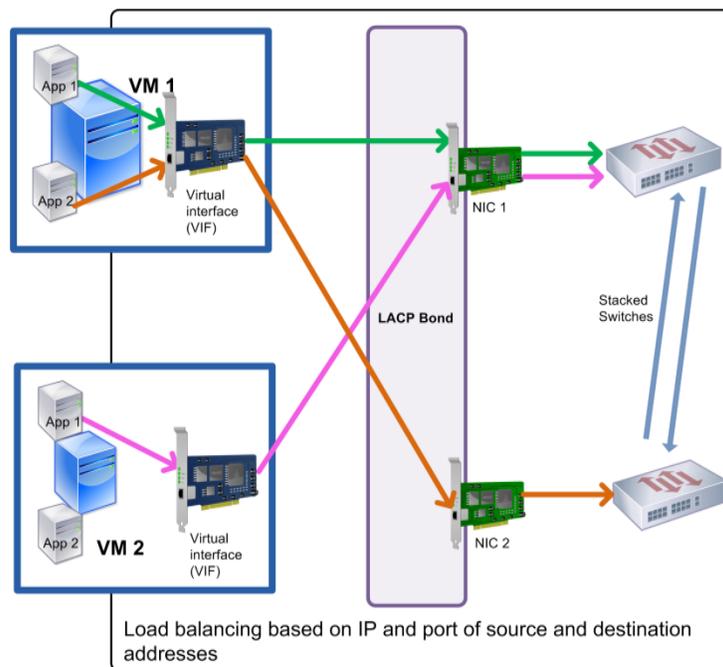
Note: You configure settings for outgoing and incoming traffic separately on the host and the switch: the configuration does not have to match on both sides.

Load balancing based on IP and port of source and destination addresses.

This hashing type is the default LACP bonding hashing algorithm. Traffic coming from one guest can be distributed over two links provided that there is a variation in the source or destination IP or port numbers.

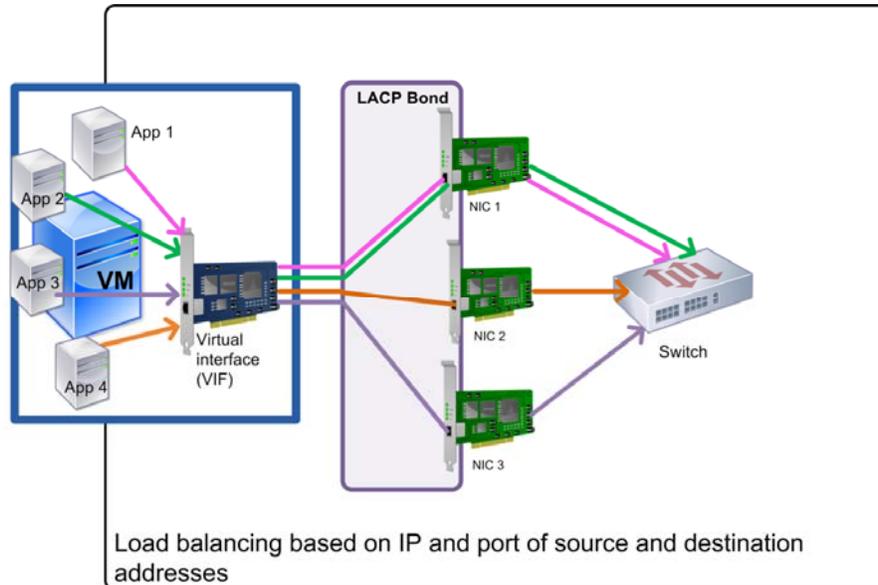
If one virtual machine is running several applications, which use different IP or port numbers, this hashing type distributes traffic over several links giving the guest the possibility of using the aggregate throughput. This hashing type lets one guest use the whole throughput of multiple NICs.

Likewise, as shown in the illustration that follows, this hashing type can distribute the traffic of two different applications on a virtual machine to two different NICs.



*This illustration shows how, if you use LACP bonding and enable **LACP with load balancing based on IP and port of source and destination** as the hashing type, the traffic coming from two different applications on VM1 can be distributed to two NICs*

Configuring LACP bonding based on IP and port of source and destination address is beneficial when you want to balance the traffic of two different applications on the same VM (for example, when only one virtual machine is configured to use a bond of three NICs).

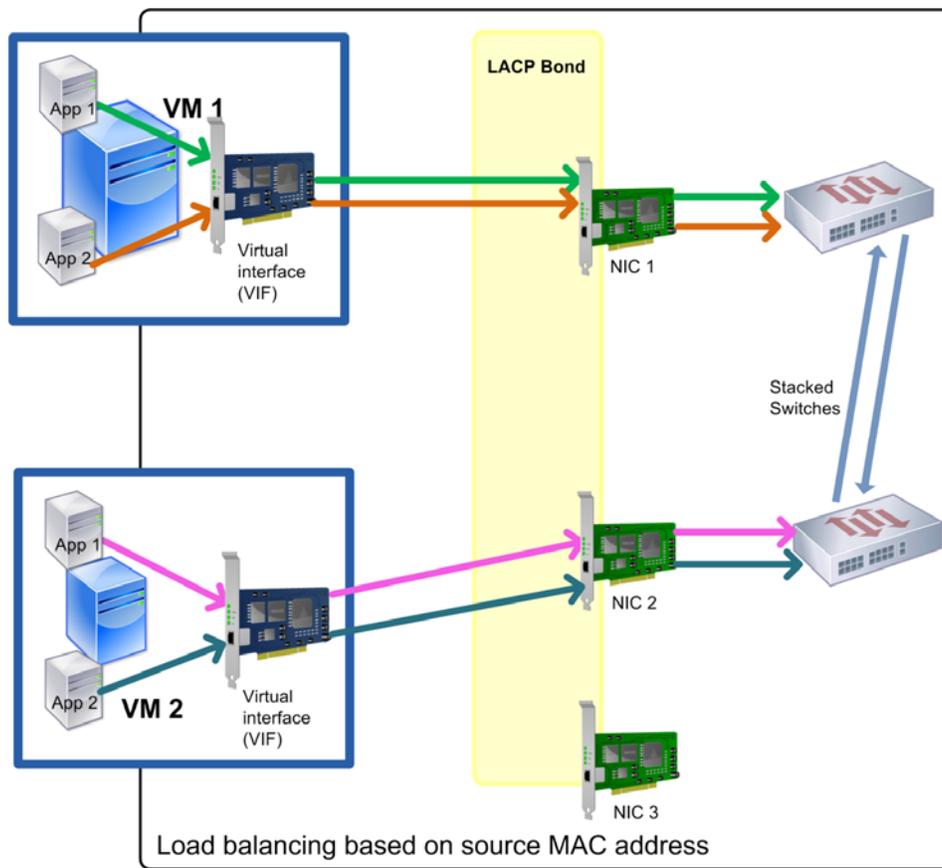


*This illustration shows how, if you use LACP bonding and enable **LACP with load balancing based on IP and port of source and destination** as the hashing type, XenServer can send the traffic of each application in the virtual machine through one of the three NICs in the bond even though the number of NICs exceeds the number of virtual network interface.*

The balancing algorithm for this hashing type uses five factors to spread traffic across the NICs: the source IP address, source port number, destination IP address, destination port number, and source MAC address.

Load balancing based on source MAC address.

This type of load balancing works well when there are multiple virtual machines on the same host. Traffic is balanced based on the virtual MAC address of the VM from which the traffic originated. XenServer sends outgoing traffic using the same algorithm as it does in the case of active-active bonding. Traffic coming from the same guest is not split over multiple NICs. As a result, this hashing type is not suitable if there are fewer virtual network interfaces than NICs: load balancing is not optimal because the traffic cannot be split across NICs.



This illustration shows how, if you use LACP bonding and enable LACP based on source MAC address as the hashing type, if the number of NICs exceeds the number of virtual network interfaces, not all NICs will be used. Because there are three NICs and only two VMs, only two NICs can be used at the same time and the maximum bond throughput cannot be achieved. The packets from one VM cannot be split across multiple VMs.

Switch Configuration

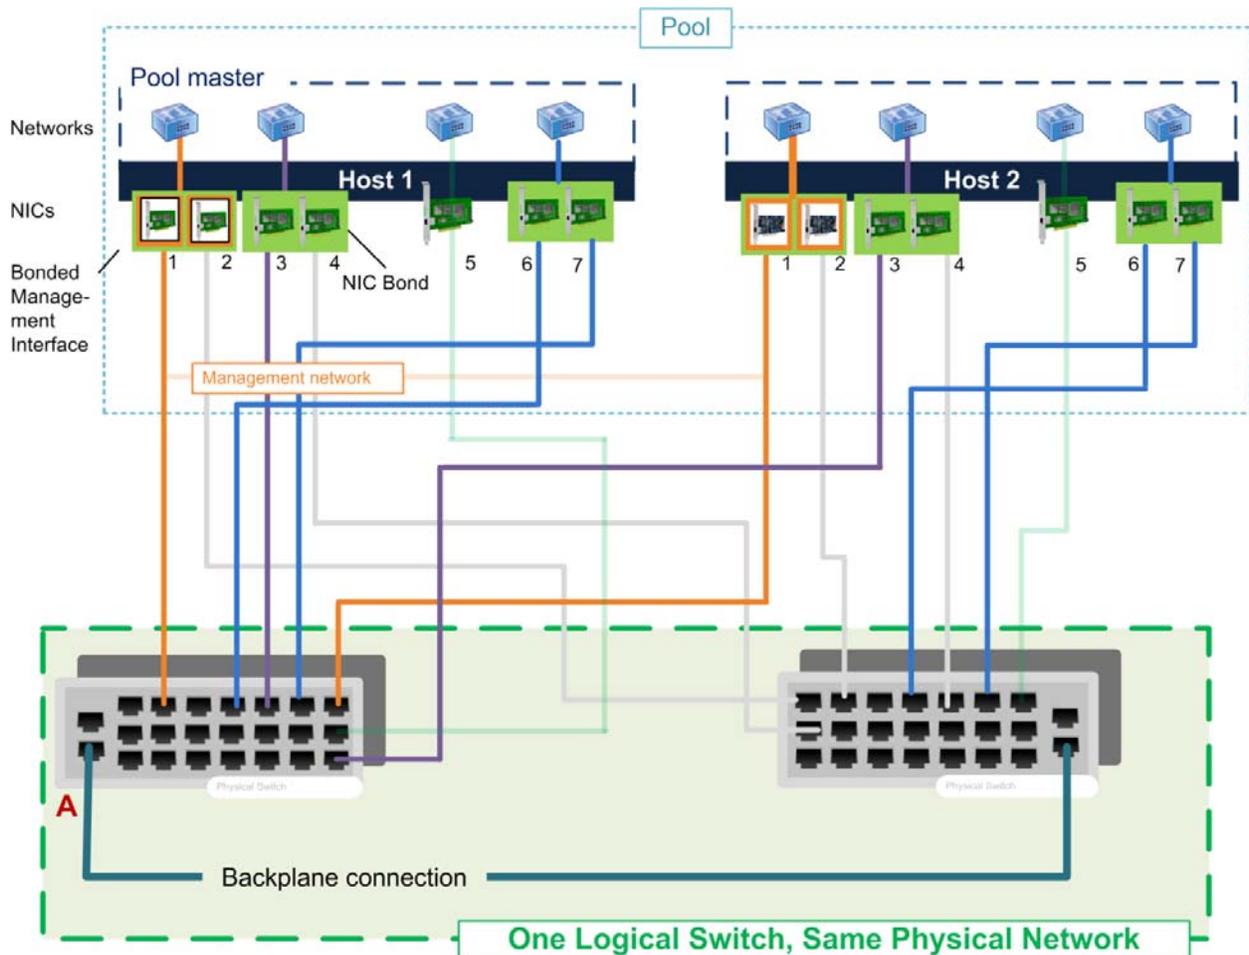
Depending on your redundancy requirements, you can connect the NICs in the bond to either the same or separate stacked switches. If you connect one of the NICs to a second, redundant switch and a NIC or switch fails, traffic fails over to the other NIC. Adding a second switch prevents a single point-of-failure in your configuration in the following ways:

- When you connect one of the links in a bonded management interface to a second switch, if the switch fails, the management network still remains online and the hosts can still communicate with each other.
- If you connect a link (for any traffic type) to a second switch and the NIC or switch fails, the virtual machines remain on the network since their traffic fails over to the other NIC/switch.

When you want to connect bonded NICs to multiple switches and you configured the LACP bonding mode, you must use stacked switches. The term *stacked switches* refers to configuring

multiple physical switches to function as a single logical switch. You must join the switches together physically and through the switch-management software so the switches function as a single logical switching unit, as per the switch manufacturer's guidelines. Typically, switch stacking is only available through proprietary extensions and switch vendors may market this functionality under different terms.

If you experience issues with active-active bonds, the use of stacked switches might be necessary. Active-passive bonds do not require stacked switches.



This illustration shows how two NICs in a bonded pair use the same network settings, as represented by the networks in each host. The NICs in the bonds connect to different switches for redundancy.

Bonding Managed Interfaces and MAC Addressing

Because bonds function as one logical unit, all NICs only have one MAC address amongst them. That is, unless otherwise specified, the bonded NICs uses the MAC address of the first NIC in the bond.

You can determine the first NIC in the bond as follows:



- In XenCenter, the first NIC in the bond is the NIC assigned the lowest number. For example, for a bonded NIC named “Bond 2+3,” the first NIC in the bond is NIC 2.
- When creating a bond using the `xe bond-create` command, the first PIF listed in the `piif-uuids` parameter is the first NIC in the bond.

When creating a bond, make sure that the IP address of the management interface before and after creating the bond is the same. When using DHCP, make sure that the MAC address of the management interface before creating the bond (that is, the address of one of the NICs) is the same as the MAC of the bond after it is created.

From XenServer 6.0 onwards, if the MAC address is not specified (using the `mac` parameter when bonding from the CLI), the bond uses the MAC address of the management interface if the management interface is one of the interfaces in the bond. If a secondary interface is in the bond (and the management interface is not), the bond uses the MAC address and IP address from that secondary interface. Otherwise, if none of the NICs in the bond is a managed interface, the bond uses the MAC address of the first named NIC.

Note: After a pool is up and running, Citrix recommends using caution when bonding the management interface.

For more information about the `xe bond-create` command, see the *XenServer Administrator's Guide*.

Best Practices for Bonded Interfaces

Citrix strongly recommends configuring NIC bonding on the pool master after you join all member servers to the pool and before you create VMs. While it is technically possible to change NIC bonding afterwards, it can create issues.

Important: Do not join a host that already has a bond configured on it to a pool without first deleting the bond.

Ideally, for maximum performance, configure NIC bonds for VM traffic and isolate management traffic on its own network. However, note the following:

- Separating management, storage and VM traffic across different NICs helps prevent contention for network resources between the management, VM, and storage networks. However, unless you bond interfaces, they do not provide redundancy.
- Using the same bonded interface rather than two separate NICs for management and storage traffic can decrease performance, as all traffic will go through the same NIC.
- Always create bonds before creating virtual network interfaces on VMs.

Warning: Do not attempt to bond NICs while the High Availability feature is enabled. Creating bonds can interrupt the in-progress High Availability heartbeat and cause hosts to self-fence



(shut themselves down). Consequently, the hosts probably will not reboot correctly and will need the **host-emergency-ha-disable** command to recover.

Both NICs in the bond must have the same frame size (MTU). You can edit the frame size in XenCenter when you create the bond. For additional information, see “Configuring Networks with Jumbo Frames” on page 85.

Network card failures are rarely the reason network connections fail. Switch failures, network outages, and performance issues on one subnet are more common. Consider making bonded interfaces more reliable by:

- Configuring two or more different network up links (subnets) for each NIC in the bond. Not only is this a requirement, it helps ensure a network connection if there are issues (unrelated to hardware or switches) on one subnet.
- Connecting NICs bonded for management, storage, and guest traffic to different redundant switches.
- In an active-passive configuration, Citrix recommends connecting the passive NICs in each bond into one switch and the active NICs in each bond into a separate switch. If one of the switches fails, you still do not have a single point of failure because the failover NIC goes into another switch.

Where bonding requirements vary by workload type, consider grouping workloads with matching bonding requirements and NIC configurations in the same pool. Citrix makes this suggestion because XenServer automatically configures identical networking configurations across all hosts on the pool.

Designing Networks for Performance

If performance is a goal for your XenServer networking configuration, consider your I/O and performance requirements.

The I/O requirements of the workloads you are virtualizing determine how you should configure the physical NICs in each server (and, by extension, each pool).

Analyzing the traffic levels of the workloads may show, for example, that traffic levels let some hosts share a common external network while other workloads may require access to dedicated NICs to support their requirements. To improve throughput, you can implement NIC bonding: XenServer uses Source Load Balancing (SLB) to share load across bonded NICs.

It should be noted, however, that bonding has the *potential* to increase throughput only when VM traffic is not balanced between NICs. If VM traffic is already balanced between separate NICs, bonding will not increase throughput.



One of the most important factors for network performance is CPU utilization. The CPU utilization of your workloads has a significant impact on network throughput. As the CPU demands of workloads increase, the effective network performance may degrade. The impact of CPU utilization is discussed throughout this section.

Review the negative and beneficial impact sections that follow for ideas of ways you can optimize the performance of your XenServer network configuration.

Negative Impact

- **Load on NICs.** Some models of network cards require firmware upgrades from the vendor to work reliably under load, or when certain optimizations are turned on. If you are seeing corrupted traffic to VMs, you should first try to obtain the latest recommended firmware from your vendor and apply a BIOS update. If this does not resolve your issue, contact Citrix Technical Support.

Potentially Beneficial

- **Improving the physical CPU.** In Citrix network throughput testing, all network throughput tests were, in the end, limited by virtual CPU capacity. Ultimately, this means that hosts with better physical CPUs typically achieve higher network throughput for the Control Domain and VMs.
- **NIC Bonding.** By bonding two or more NICs together, you can increase the total amount of throughput available by better balancing the traffic. While bonding NICs does not affect the total bandwidth available—the bandwidth available is the same as the combined bandwidth of the individual NICs – bonding can make better use of available resources.
- **Upgrading the NIC.** Provided infrastructure that can support such interfaces is in place, upgrading the NIC (for example, from a 1 gigabit NIC to a 10 gigabit NIC can improve performance).
- **Implementing jumbo frame** support can improve performance for storage traffic. For more information about using jumbo frames with storage, see “Chapter 6: Designing Your Storage Network Configuration” on page 76. Enabling jumbo frames can also improve network performance for VM-to-VM traffic (when the VMs are on different hosts).
- **Implementing SR-IOV** for Provisioning Services server traffic. As of XenServer 5.6 Feature Pack 1, XenServer supports SR-IOV. For more information, see “Virtualizing the Provisioning Services Server” on page 88.
- **Tweaking TCP settings inside the VMs.** According to Citrix testing, tweaking TCP settings inside the VMs can lead to substantial network performance gains. The main reason for this is that most systems are still, by default, configured to work well on 100Mb/s or

1Gb/s NICs--and not 10Gb/s NICs. You can find the settings Citrix recommends in the Xen.org [Network Throughput Guide](#).

Resolving Network Bottlenecks

There are many potential causes for networking bottlenecks. The following list of symptoms provides likely causes and guidance:

- **I/O is extremely slow on the virtual machine.**
 - **Verifying the symptom:** Compare the results of an I/O speed test on the slow VM to a healthy VM; the healthy VM is ten times faster than the slow VM.
 - **Probable cause:** The slow VM does not have paravirtualized drivers installed.
 - **Background:** With paravirtualized drivers, a VM can make direct use of some of the underlying hardware, leading to better performance.
 - **Recommendation:** Install paravirtualized drivers.

- **VM's VCPU is fully utilized.**
 - **Verifying the symptom:** Run `xentop` in the Control domain. This should give a fairly good estimate of aggregate usage for all the VCPUs in a VM: pressing `V` reveals how many seconds were spent in which VM's VCPU. Running VCPU measurement tools inside the VM *does not* give reliable results; these tools can only be used to find rough relative usage between applications in a VM.
 - **Background:** When a VM sends or receives network traffic, it needs to do some basic packet processing.
 - **Probable cause:** There is too much traffic for that VCPU to handle.
 - **Recommendation 1:** Try running the application that sends/receives the network traffic with multiple threads. This gives the operating systems a chance to distribute the workload over all available VCPUs.
 - **Recommendation 2:** Try enabling NIC offloading --- as described in the in the Xen.org [Network Throughput Guide](#) – however, it should be noted that NIC offloading is not fully supported.

- The VM's first (and possibly only) VCPU is fully utilized.
 - **Verifying the symptom:** See the method for the previous symptom.
 - **Background:** Currently, only a VM's first VCPU can process the handling of interrupt requests.

- **Probable cause:** The VM is receiving too many packets for its current setup.
- **Recommendation 1:** If the VM has multiple VCPUs, try to associate application processing with non-first VCPUs.
- **Recommendation 2:** Use more (1 VCPU) VMs to handle receive traffic, and a workload balancer in front of them.
- **Recommendation 3:** If the VM has multiple VCPUs and there is no definite need for it to have multiple VCPUs, create multiple 1-VCPU VMs instead (see Recommendation 2).

Testing XenServer Network Performance

You can observe a host's network performance in several ways, including using the XenCenter **Performance** tab, Xen commands like **xentop** and **xenmon**, and networking testing tools, such as Iperf.

The XenCenter **Performance** tab displays the number of Mbps each NIC sends and receives and the load on the Control Domain, and the utilization for each CPU on the host. By clicking **Configure Graphs**, you can also change the Data Source to see network send and receive errors.

Before you begin testing the performance of your XenServer network configuration, you should understand what performance your hardware can theoretically support.

To test the performance of your XenServer network configuration on a specific XenServer host:

1. Start by using a network testing tool (such as Iperf).
2. Examine the throughput of your network, as your network testing tool recorded it. Compare the recorded throughput level with what your hardware can theoretically support. For example, if you know that your NIC supports 1 gigabit of throughput, then you should be getting a number close to that as your network throughput keeping in mind it is not possible to achieve the limit since there is always some overhead.
3. If you are not getting this level of throughput, do one or more of the following to find out more information:
 - In XenCenter, select the host, and use the **Performance** tab in XenCenter to check the following:

Make sure that VMs do not have a CPU, memory, or disk bottleneck because this prevents the VMs from receiving full throughput. The most common cause of reduced throughput is high CPU utilization by guests.



- Run the **xentop** command. The **xentop** command shows you how much CPU each VM uses, network traffic to and from the VM, and disk traffic on the Control Domain. For more information about this command, enter **xentop -help** at the command prompt on your host.
- Run the **xenmon** command. The **xenmon** command can help identify which domains and VMs are creating the highest I/O or processing loads on a host. For more information about this command, enter **xenmon.py -help** at the command prompt on your host.

What to do next?

1. Experiment with the number of processing threads (vs. the available virtual CPUs) on the VMs.
2. Reconsider your VM to host ratio. You may need to reduce the number of VMs on a host, or host a different mixture of workloads, to obtain the network performance you want.
3. Make sure there are no other VMs sending or receiving a lot of network traffic. Consider changing the home servers VMs are assigned to so they are balanced across different hosts according to their network utilization. You may also want to consider configuring Workload Balancing to make balancing recommendations or automatically balance VMs according to reads and writes.
4. Ensure VM traffic is evenly distributed across physical CPUs. For more information, see CTX127970—[Distributing Guest Traffic Over Physical CPUs in XenServer](#).
5. Load balance your network traffic across NICs. Make sure that no one NIC has an excessive number of VMs pointed to it and/or that these VMs are not sending an excessive amount of traffic.
6. Separate the management, storage, and VM traffic, as described in throughout this guide, to see if this improves performance.
7. If you have a mixed environment, consider putting a mixture of Linux and Windows VMs on that host. When virtualized, the Linux operating system usually puts less stress on the Control Domain and CPU resources.

Citrix recommends running these tests on a pilot environment before finalizing your hardware requirements and overall design. These tests may indicate that you need additional network connections, to rearrange workload groupings, to purchase more NICs, purchase NICs that support different configurations than you originally intended (for example, jumbo frames or SR-IOV) and so on.



Limiting Bandwidth Consumption for High Demand Workloads

In environments where the workloads send a lot of data and can potentially consume a lot of network bandwidth, you might want to limit the amount of data these workloads can send and slow down the transmission rate. This helps ensure other VMs on the host receive adequate network transfer rates.

To limit data transfer speeds, you can specify a maximum transfer rate in kilobytes per second (or kilobits per second in the CLI) when you create the VM's virtual network interface(s) or at a later time. When you limit the transfer rate, you are doing so for all the data that VM sends over the virtual network interface to its associated network link. If that VM uses multiple network links and you want to limit the transfer rate for all of them, you must create a QoS limit in each virtual network interface for each network link.

Setting limits can be particularly useful when different organizations own different VMs on the same host since it helps ensure each VM gets a fair share of network bandwidth.

If you want to limit data transmission on more than one virtual network interface on a VM, you must configure each virtual network interface to have a QoS limit.

To configure a QoS limit for VM output, you have several options:

- **In XenCenter.** Select the VM, click the **Network** tab, and click either **Add Interface** or **Properties**.
- **In the vSwitch Controller.** For pools using the vSwitch and the Distributed Virtual Switching solution, configure the QoS setting in vSwitch Controller.
- **Using the vif-param-set xe command.** For example, to limit a virtual network interface to a maximum transfer rate of 100 kilobits per second (kbps), use the command:

```
xe vif-param-set uuid=<vif_uuid> qos_algorithm_type=ratelimit
xe vif-param-set uuid=<vif_uuid> qos_algorithm_params:kbps=100
```

You do not need to create a QoS limit for all virtual network interfaces/VMs that use that network link. You only need to set limits on the VMs you want to constrain.

Example: Calculating Bandwidth Limits

You can determine bandwidth limits using a formula like the following:

$$\text{Max Rate Transfer per VIF} = (\text{NIC speed in kilobits per second} / \text{number of VIFs}) * \text{desired network utilization \% (e.g., 70\%)}$$

An example of the application of this formula is as follows:



You have five VMs on a host with one virtual network interface on each VM. The virtual network interfaces all use the same 10 gigabit Ethernet NIC. You want to limit bandwidth consumption for all five VMs and their virtual network interfaces.

Assuming a desired network utilization of 70%, you can theoretically limit each virtual network interface bandwidth to 1,468,006 kilobits per second (or 183,500 kilobytes per second).

To obtain this number, we used the following calculations:

1. A 10Gigabit NIC = 10,485,760 kilobits per second
2. $10,485,760 / 5$ virtual network interfaces = 2,097,152 kilobits per second (per virtual network interface)
3. $2,097,152 * 70\%$ = 1,468,006 kilobits per second maximum transfer rate (per virtual network interface)
4. $1,468,006$ kilobits per second $/ 8 = 183,500$ kilobytes per second. XenCenter requires you enter the rate limit in kilobytes per second.

After determining the maximum transfer rate per virtual network interface, you can then increase or decrease this value on each VM according to its bandwidth needs.

The best practice is not to guess at the network utilization, but rather measure the actual throughput from the NIC and divide it by the number of VMs using that NIC. If, for example, you are achieving 8,074,035 kilobits per second divide that by the number of VMs with a virtual network interface configured to use that NIC. For example, $8,074,035 / 8$ VMs = an average throughput of 1,009,254 kilobits per second available for each VM.

You might decide to limit some VMs to far less than this value (for example, 800,000 kilobits per second) and others to more (for example, 1,200,000 kilobits per second) depending on the bandwidth requirements and business priority of the workloads.

However, be careful not to exceed the total amount of throughput achievable. If you do exceed it, the limits you set may not be reached.

In general, if you have not measured the throughput, it is better to set the maximum transfer rate slightly lower than what you expect your NICs can realistically achieve.

Note: The virtual network interface QoS rate limit setting is different than the virtual disk QoS disk priority setting. The virtual disk QoS setting is a storage configuration that lets you assign priorities to virtual disks so that disks with higher priority are accessed faster than other disks. The *XenServer Administrator's Guide* describes both of these settings in more depth.

Additional Considerations

This section discusses an additional consideration for your networking configuration.



Creating a Dedicated Network for Storage Motion

Storage XenMotion is a XenServer feature that lets you move VMs from one host to another, where the VMs are not located on storage shared between the two hosts. As a result, you can migrate VMs stored on local storage without downtime and VMs can be moved from one pool to another.

Storage XenMotion requires you specify a NIC with an IP address configured on it as the network for storage migration. By default, Storage XenMotion uses the management interface as its migration network.

Depending on the size of virtual machines and virtual disks you want to migrate, you may want to dedicate a NIC and network for storage migration. Configuring a different NIC for storage motion prevents storage migration traffic from competing with management traffic. If you use the management interface for storage migration and perform storage migration frequently or have large virtual disks, it can use a lot of bandwidth and impact the performance of their pool (by reducing bandwidth for management traffic).

For more information about Storage XenMotion and its requirements, see the *XenServer 6.1.0 Virtual Machine User's Guide*.

Enabling Switch-port Locking

The XenServer switch-port locking feature lets you control traffic being sent from unknown, untrusted, or potentially hostile VMs by limiting their ability to pretend they have a MAC or IP address that was not assigned to them. You can use the port-locking commands in this feature to block all traffic on a network by default or define specific IP addresses from which an individual VM is allowed to send traffic. For more information, see the *XenServer Administrator's Guide*.

Enabling Promiscuous Mode for Traffic Monitoring

XenServer supports promiscuous mode, an Ethernet configuration for NICs in which the NIC receives all traffic on its link instead of only the frames addressed to that NIC's MAC address. Organizations may use promiscuous mode for a variety of reasons, such as requirements from transparent proxies or specialized traffic monitoring, security, and troubleshooting applications.

You can enable promiscuous mode at both the VM and XenServer host (physical server) levels. That is, you can configure promiscuous mode for both the virtual network interfaces and physical NICs.

When you enable promiscuous mode on virtual network interface or physical NIC, the mode lets you see all the traffic on a virtual switch. You might want to enable promiscuous mode when you want to:

- Run software (for example, software for a switch) or integrate an appliance that requires visibility into all traffic passing across the physical NIC to which it is connected. For configuration instructions, see CTX116493—[How to Enable Promiscuous Mode on a Physical Network Card](#).

- See all traffic going across the network (specifically, the virtual switch) between the NIC (PIF) and a VM's virtual interface. Instructions for implementing this configuration appear in CTX121729—[How to Configure a Promiscuous Virtual Machine in XenServer](#).

If your goal is to see all traffic VMs send across a specific network across a pool, you may need to configure promiscuous mode on a virtual network interface on a VM in every host in the pool.

When determining if you want to enable promiscuous mode, consider the CPU load on the XenServer host. When XenServer runs in promiscuous mode, the kernel receives all network traffic and the CPU utilization on the host increases. Because this can slow down responses to network packets, this in turn can also increase network latency.

If you are using the vSwitch Controller introduced in XenServer 5.6 Feature Pack 1, you can use RSPAN instead of promiscuous mode to display traffic. The vSwitch Controller also includes functionality for mirroring.

Promiscuous Mode and Security

When promiscuous mode is configured on a virtual or physical switch port, the virtual network interface or NIC connected to this switch port receives all traffic that travels through the switch and then passes it to the VM that “owns” the virtual network interface (or, in the case of NICs, the Control Domain). This means that if malware or another type of malicious attack reaches across your network all VMs will become infected simultaneously.

Citrix recommends that you limit your use of promiscuous mode to troubleshooting and, possibly, security monitoring.

Warning: Do not enable promiscuous mode without a good reason since the VM in promiscuous mode can access the network traffic for other VMs.

Chapter 6: Designing Your Storage Network Configuration

This section is intended to highlight storage considerations you should include while designing your XenServer network configuration. It includes the following topics:

- The need to create a separate storage network
- How to assign IP addresses to NICs
- How to improve performance for storage traffic
- Guidance about choosing between NIC bonding and multipathing for storage

An extensive discussion of storage configurations is beyond the scope of this document.

Overview

This chapter is designed to help you configure your IP-based storage configuration so that it has redundancy and gets the best network performance with the lowest impact on other network traffic.

Specifically, this chapter explains how to create a separate storage network and when you would want to do so. It also explains the differences between the two failover choices for storage traffic, bonding and iSCSI multipathing. It also provides information about assigning IP addresses to NICs, since this is a common requirement for IP-based storage.

Other techniques discussed in this chapter that improve performance include jumbo frames.

Connecting to Storage Controllers

To connect a storage controller to a XenServer NIC, the NIC must have an IP address. While you can re-use the management interface, the best practice is to configure a separate NIC for storage and assign it an IP address. A NIC with an IP address is known as a *secondary interface*. When you create a secondary interface for storage, you are defining an IP address for a specific NIC on the XenServer host. This IP address will be the address the storage uses to communicate with XenServer.

If you want to create two secondary interfaces for storage (that is, have two different networks for storage), each secondary interface must be on a separate subnet. For example, if you want to configure two secondary interfaces for storage, XenServer requires IP addresses on three different subnets – one subnet for the management interface, one subnet for Secondary Interface 1, and one subnet for Secondary Interface 2.

As described previously, when you create a secondary interface for storage, you are not necessarily creating a dedicated storage network unless you:

- Ensure the physical configuration segregates traffic
- Configure VMs with only guest networks
- Ensure network forwarding is turned off

Also, every host in your pool must have a parallel matching networking configuration. This means that if you configured NIC2 as Secondary Interface 1 on the 10.204.150.x subnet, then you must configure every NIC2 on each host in your pool to use that same subnet and carry storage traffic to the same controller.

Creating a Separate Storage Network

Citrix recommends dedicating one or more NICs as a separate storage network for NFS and iSCSI storage implementations. Many consider creating a separate storage network to be a best practice.

By configuring *secondary interfaces*, you can both assign an IP address to a NIC and isolate storage and network traffic, provided the appropriate physical configuration is in place. The term *managed interface* refers to any NIC assigned an IP address for identification purposes. (The management interface, which is introduced in “Networking Configuration after Installation” on page 14, is a *type* of managed interface.)

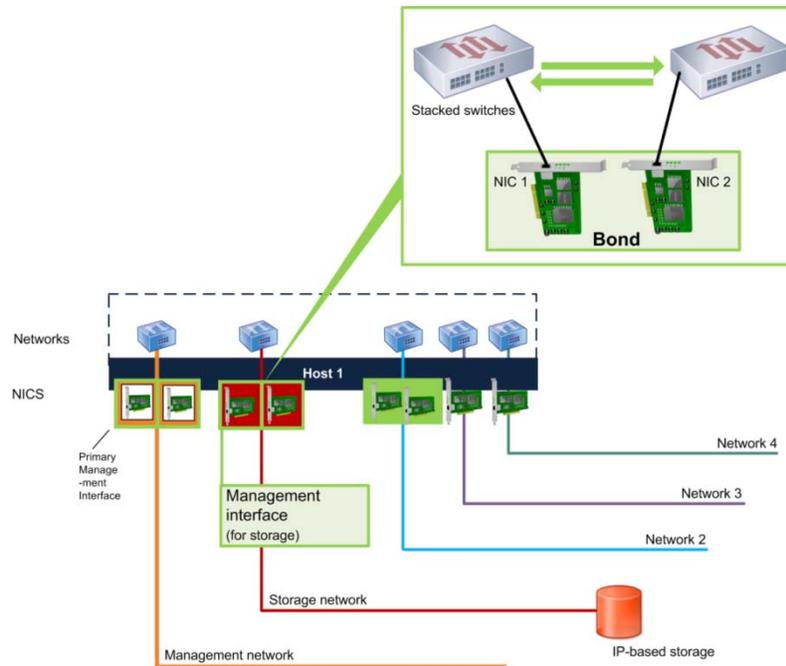
Tip: To figure out which NIC is the management interface, in XenCenter, click the **Network** tab, in the **IP Address Configuration** section, check the **Interface** column for the word “Management.”

You can segregate traffic by configuring a secondary interface(s) for storage and configure XenServer to access storage through that interface. Then, physically isolate the guest network and configure virtual machines to only use that isolated network. Ideally, the network should be bonded or use multipathing, as described in “Choosing to Enable Multipathing Support or Bond NICs.”

The overall process for creating a separate storage network is as follows:

1. Configuring physical network infrastructure so that different traffic is on different subnets.
3. Creating a secondary interface to use the new network.
3. Configuring redundancy, either multipathing or bonding.

In the illustration that follows, an administrator created a NIC bond from NICs 3 and 4 and then configured the bonded pair as a secondary interface for storage.



*This illustration shows how the bond made from NICs 3 and 4 is configured as a secondary interface. XenServer sends its storage traffic over this NIC bond onto the storage network and, ultimately, the storage array. The exploded diagram shows how each NIC in the bond connects to a different switch. Note: The need for stacked switches varies according to the bond mode. See page **Error! Bookmark not defined.***

Creating secondary interfaces lets you establish separate networks for, for example, IP-based traffic provided:

- You do not configure XenServer to use this network for any other purpose (for example, by pointing a virtual network interface to this network).
- The appropriate physical network configuration is in place.

For example, to dedicate a NIC to storage traffic, the NIC, storage target, switch, and/or VLAN must be configured (physically connected) so that the target is only accessible over the assigned NIC.



To ensure that the storage traffic is separated from the management traffic, the storage network must be on a different subnet. The subnet for storage must be a separate IP subnet that is not “routable” from the management interface. If the physical or logical configuration does not enforce the traffic separation, then XenServer may direct storage traffic over the management interface after a host reboot, due to the order in which XenServer initializes NICs.

In smaller environments, routing guest, management, and/or storage traffic together may not matter. However, in larger environments or environments with strict security policies, you may want to separate your traffic.

In order for IP-based storage, such as iSCSI software initiator and NFS, to communicate with XenServer, the XenServer NICs must have IP addresses. To specify an IP address for a NIC or bond, you must create a secondary interface or reuse the IP address assigned to the management interface.

In other words, you can assign a XenServer IP address for your storage array to connect to by either:

- Configuring secondary interfaces so you can assign IP addresses to NICs if you do not want to use the management interface for routing storage traffic.
- Routing storage traffic through the management interface—since the management interface has an IP address this will work.

Some environments, such as test labs or at very small companies, may experience little impact from routing management and storage traffic on one interface. However, in general, Citrix strongly recommends that you do not route storage traffic over the management interface.

If you want to bond the management interface, create the management interface first and then bond it. Ideally, the first NIC in the bond should be the one for that has the IP address assigned to it. However, Citrix also recommends configuring the MAC address on the bond if you are using DHCP, as described in “Bonding Managed Interfaces and MAC Addressing” on page 65.

Important: Each managed interface, including the management interface, must be on a separate subnet. For example, if you want to configure IP addresses on two more NICs for storage, you will require IP addresses on three different subnets – one subnet for the management interface, one subnet for Secondary Interface 1, and one subnet for Secondary Interface 2.

Assigning IP Addresses to NICs (Secondary Interfaces)

You can assign IP addresses to NICs by configuring secondary interfaces in XenCenter and using the `xe` commands. To do so in XenCenter:

1. Ensure that the NIC is on a separate subnet, or routing is configured to suit your network topology in order to force the desired traffic over the selected NIC.

- In the XenCenter resource pane, select the host that is the pool master, click the **Network** tab.

Pool Networks

Networks

Name	Description	NIC	VLAN	Auto	Link Status	MAC	MTU
Bond 1+2		Bond 1+2	-	No	Connected	00:22:19:53:18:9c	1500
Network 0		NIC 0	-	Yes	Connected	00:22:19:53:18:9a	1500
Network 3		NIC 3	-	Yes	Connected	00:15:17:97:bf:d5	1500

Add Network... Properties Remove

IP Address Configuration

Server	Interface	Network	NIC	IP Setup	IP Address	Subnet mask	Gateway	DNS
host-34	Management	Network 0	NIC 0	Static	10.204.154.34	255.255.255.0	10.204.154.1	10.204.154.177
host-35	Management	Network 0	NIC 0	Static	10.204.154.35	255.255.255.0	10.204.154.1	10.204.154.177

Configure...

- The **IP Address Configuration** feature is found on this tab. Instructions for configuring management interfaces varies by XenCenter release. See the *XenCenter Help* for instructions for your release.

Configure IP Addresses - host-35

Configure IP address settings for server management traffic (the management interface) on 'host-35' here. You can also assign IP addresses for dedicated storage or other types of network traffic (secondary interfaces).

Management
Network 0; Static

Storage 1
<None>

Add IP address

Storage 1

Name: Storage 1

Network: Bond 1+2

IP address settings:

Automatically obtain settings using DHCP

Use these settings:

IP address:

Subnet mask:

Gateway:

Remove IP Address

[Tell me more about configuring IP-based storage traffic](#)

OK Cancel

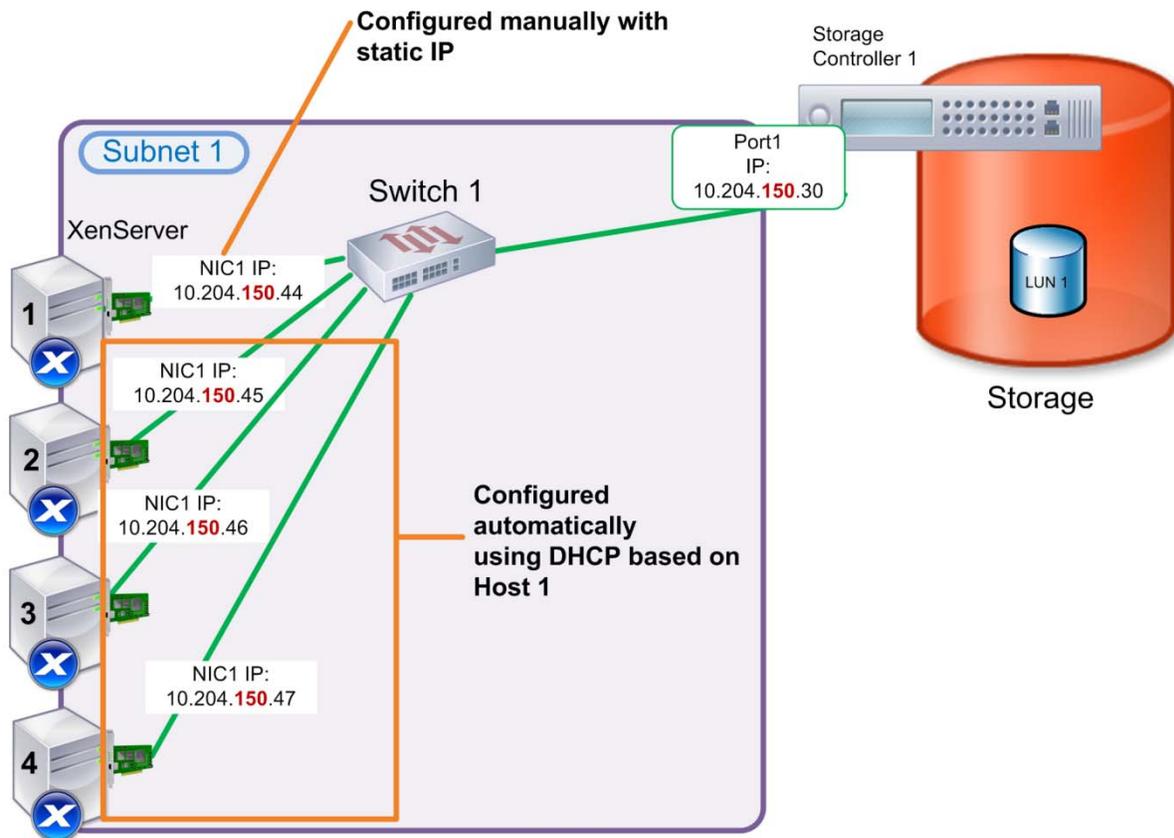
Tip: If you want to dedicate a NIC for storage, check XenCenter to make sure that the NIC's associated network is not configured so that it is added to the VMs by default. To do so, in the **Networks** tab, right-click *<your-storage-network>* > **Properties**. Click **Network Settings** and

make sure the **Automatically add this network to new virtual machines** check box is not selected.

Assigning IP Addresses on All Hosts in a Pool

A simple way to assign IP addresses to NICs across a pool is to assign the NIC on the first host in the pool a static IP address and then let DHCP assign the corresponding NICs on the other hosts their IP addresses.

After you specify the static IP address, XenServer automatically creates a secondary interface on the corresponding NICs on each host and sets them to use the next available IP address, as shown in the illustration that follows.



This illustration shows how after an administrator created a secondary interface on NIC 3 on Host 1 and set it to 10.204.150.44, XenServer automatically created secondary interfaces on the other hosts in the pool and set their IP addresses using DHCP. Each host in the pool is assigned the next IP address in the range –Host 2 is set to 10.204.150.45, Host 3 is set to 10.204.150.46, and Host 4 is set to 10.204.150.47.



Important: If you are configuring IP addresses on NICs to connect to an iSCSI storage controller, the static IP address you enter must be in the same subnet as the storage controller to which you will ultimately be connecting the NIC.

The following procedure explains this technique.

To assign all corresponding NICs in a pool an IP address

1. Ensure that the NICs are on a separate subnet, or routing is configured to suit your network topology in order to force the desired traffic over the NICs.
2. In the **Resource** pane of XenCenter, click on the pool. Click the **Network** tab, and then click the **Configure** button.
3. In the left pane, click **Add IP address**.
4. Give the NIC a meaningful name, and select the Network associated with the NIC where you want to assign the IP address. An example of a meaningful name might be “Storage Network 1” or “*your-storage-controller-name-or-model 2.*” Numbering your storage management interfaces will make it easier to configure multipathing, if desired, later on.
If you want to assign the IP address on NIC 3, for example, select the network associated with NIC 3.
5. Select **Use these network settings**, and enter a static IP address you want to configure on the first NIC, the subnet mask, and gateway, and click OK.

The range of IP addresses XenServer will use to configure the corresponding IP addresses in the pool appears to the right of the IP address you entered.

After you click **OK**, XenServer assigns the IP addresses to the NICs. When you click the Pool node in the Resource pane, the new secondary interfaces appear in the **IP Address Configuration** section of the **Networking** tab.

Configuring Redundancy for Storage Traffic

For environments that want redundancy for their network storage traffic, XenServer supports NIC bonding and multipathing, including iSCSI HBA multipathing and software iSCSI multipathing. The term *multipathing* refers to routing storage traffic to a storage device over multiple paths for redundancy (failover) and increased throughput.

This section provides information about when to choose NIC bonding instead of multipathing and how to configure iSCSI multipathing.

Bonding NICs or configuring multipathing helps provide redundancy in case of partial network failure. Redundancy helps prevent interruptions to disk reads and writes. Interrupting disk reads and writes can lead to guest operating system failure if the VM consequently disconnects from the remote disk.



Citrix strongly recommends that you do not mix NIC bonding and iSCSI multipathing. There is no benefit from layering multipathing and NIC bonding on the same connection. After you enable multipathing, you not only have better performance but you also have the failover that bonding would have provided.

Note: XenServer supports multipathing for Fibre Channel and iSCSI SANs. However, multipathing for Fibre Channel SANs is not discussed in depth in this guide since traffic for Fibre Channel goes across a Fibre Channel network and not a TCP/IP network. For information about Fibre Channel multipathing, see the *XenServer Administrator's Guide* and the Citrix Knowledge Center.

Choosing to Enable Multipathing Support or Bond NICs

Citrix recommends configuring multipathing instead of NIC bonding whenever possible.

Like NIC bonding, multipathing provides failover for storage traffic. However, unlike NIC bonding, XenServer can send traffic down both paths when you configure multipathing: multipathing is an active-active configuration. By default, multipathing uses round-robin mode load balancing, so both routes will have active traffic on them during normal operation, which results in increased throughput.

Citrix recommends using multipathing when you have block-based throughput (for example, iSCSI software initiator traffic). The exception to this is if, when the storage array is connected to XenServer, it does not work with multipathing.

Even though LACP bonding provides load balancing for storage traffic, Citrix still recommends configuring multipathing whenever possible for the following reasons:

- Multipathing can extract the path priorities from the array, which is valuable for redundant controller configurations such as ALUA.
- Some arrays must be sent specific commands to reconfigure the controllers to cause a failover. (In the `multipath.conf` file, the `hardware_handler` setting indicates how to do this.)
- You can configure multipathing to perform round-robin within a priority group, with a configurable number of I/O requests on each path before going to the next. Some arrays require this configuration to coalesce requests for performance reasons. LACP bonding may send all the traffic down a single link because LACP bonding relies on different endpoints and the endpoints are fixed in this case.

The likelihood of LACP bonding matching the performance and load-balancing of multipathing for storage traffic varies on an array-by-array basis. For the previously mentioned reasons, Citrix suggests considering a multipathing configuration instead of LACP bonding for storage traffic whenever possible.

Consider using NIC bonding instead of multipathing when:

- You have an NFS storage device

- Your storage device does not support iSCSI connections over multiple IPs (for example, Dell EqualLogic or HP LeftHand SAN)

If you want to use NIC bonding for storage traffic, consider configuring the vSwitch as your networking stack and configuring LACP as the bonding type. LACP bonding enables XenServer to send storage traffic across both links in the bond. However, unlike the Linux bridge NIC bonding, LACP only support two NICs in the bond. For more information, see *Considering NIC Bonding* on page 56.

The following table shows the supported protocols for multipathing and NIC bonding:

	Multipathing	NIC Bonding
Supported Storage Protocols	Fibre Channel, iSCSI HBA, iSCSI software Initiator	NFS, CIFS, iSCSI software Initiator

Tip: To determine what XenServer hosts have multipathing enabled on them, check **Multipathing** in the **General** tab in XenCenter.

For information about configuring iSCSI multipathing, see CTX129309—[Configuring iSCSI Multipathing Support for XenServer](#).

Suggestions for Improving Storage Network Performance

This section provides guidance for improving storage network performance for iSCSI storage and jumbo frames.

Other methods of improving storage network performance were discussed in this chapter. These include creating a separate physical network for storage traffic, and when possible, choosing multipathing for redundancy since it both links are active, which improves throughput over NIC bonding.

iSCSI Storage

Provided iSCSI storage is configured correctly it can provide performance comparable to Fibre Channel.

iSCSI SANs consumes more CPU cycles on the host than NFS and Fibre Channel do. Because CPU utilization can affect network performance, this should be considered when sizing hardware and designing pools. However, it should also be noted that using an iSCSI host bus adaptor (HBA) can offload the processing for higher performance.



Citrix recommends using high performance network switches for iSCSI SANs to achieve better iSCSI performance. Citrix also recommends using redundant iSCSI network connectivity for all implementations.

It is particularly important to separate iSCSI storage traffic from management and VM guest traffic since it can interfere with non-storage traffic.

Configuring Networks with Jumbo Frames

Configuring XenServer networks to use jumbo frames can improve performance for storage traffic. Jumbo frames are Ethernet frames containing more than 1500 bytes of payload. Jumbo frames are typically used to achieve better throughput, reducing the load on system bus memory, and reducing the CPU overhead.

Currently, XenServer only supports jumbo frames if the vSwitch is used as the networking bridge on all hosts in the pool.

When determining whether your network performance will benefit from jumbo frames, consider the following:

- Jumbo frames can help offload CPU overhead.
 - By increasing the Ethernet frame size to 9000 bytes, jumbo frames reduce the number of packet headers the CPU must process, which consequently decreases the demands on the CPU. Jumbo frames also reduce the number of NIC interrupts needed when transmitting multi-packet file transfers.
 - Jumbo frames may help with slower CPUs if your NICs do not have a TCP Offload Engine (TOE) support. Jumbo frames are less apt to reduce CPU overhead with more intelligent gigabit NIC cards since these cards can process more of the packet headers on their own.
- Different types of workloads may have their own optimum packet sizes.
 - For example, for storage traffic, in general, throughput is more important than latency. As a result, if you configure larger packets for storage traffic (through jumbo frames), the storage traffic may not be negatively affected by latency and may benefit from the increased throughput.
 - For transactions that require high response times, enabling jumbo frames may not be helpful. Jumbo frames may require more buffering in the network, and, as a result, the latency for a particular item of data might be higher. (Larger packets take longer to assemble and more bandwidth per packet.) As a result, if you need high response times, bigger packets are not as good as smaller packets.



Generally, if the speed of the network and the need for high throughput increases, it is probably be good to increase the size of the packets.

- When the transfer size is, on average, relatively small, the benefits from jumbo frames might not be significant.
- The performance gains from jumbo frames vary according to the type of protocol and the type of traffic. When VM traffic is on the same network as storage traffic, if the VM traffic is latency sensitive, for example Voice-Over-IP traffic, configuring Quality of Service priorities in switches and on virtual network interfaces may be necessary to ensure the performance of the voice traffic.
- Traffic that has relatively large block sizes at the application layer may benefit from jumbo frames since large block sizes make it easier for the TCP/IP stack to use large frames.

You can enable jumbo frame support on either of the following types of XenServer networks:

- External networks.
- Private networks. For example, you could enable jumbo frame support on a cross-server private network.

However, the equipment (NICs, switches) between all links must support jumbo frames. Certain types of traffic, such as IP-based storage traffic, may benefit from jumbo frame support. You can configure XenServer to use frames of between 1500 to 9216 Mbps.

When you create the network in which XenServer will transmit data in jumbo frames, you must specify the Maximum Transmission Unit (MTU) value the network supports. In XenCenter, you do so when you create the network in the **New Network** wizard by entering the value your entire network supports in the **MTU** text box.

Citrix strongly recommends specifying jumbo frames when you create the secondary interface for the storage network.

To implement jumbo frame support for bonded networks, you specify the MTU for both NICs in the bond when you create the bond. In XenCenter, there is an **MTU** text box in the **Create Bond** dialog box and in the **New Network** wizard.

Requirements for Using Jumbo Frames

- All NICs transmitting jumbo-frame traffic must support the transmission unit speed you want to configure and be on the *XenServer Hardware Compatibility List*.
- The network segment where you want to transmit the jumbo frames must have equipment (for example, the NICs and switches) that supports the frames in all segments, from end-to-end.



- When creating virtual network interfaces and configuring NICs (PIFs), you must configure all interfaces on the network where you want to support jumbo frames with the same number of MTUs.

Citrix also recommends ensuring all components on the data path are tested for interoperability.

To achieve optimum performance, Citrix recommends that networks carrying jumbo frames use equipment with the same transmission speed (for example, 1 Gigabit). This is to promote the efficiency gains you achieve in Ethernet from standardization. While 10/100 Mbps networks may support jumbo frames, for optimum performance Citrix suggests using a minimum of 1 Gigabit Ethernet equipment. Ideally, networks should not use a mixture of 10/100 Mbps equipment and 1 Gigabit Ethernet equipment due to interoperability issues.

Additional Information

If you change MTU on the network in XenCenter, it will automatically change the MTU on the virtual network interfaces and all the virtual network interfaces and NICs/networks (switches) that point to the network supporting jumbo frames on all hosts.

Chapter 7: Considering Network Performance for PVS **– XenServer Deployments**

This chapter provides techniques for ensuring network performance when using XenServer to host Provisioning Services VMs. It includes the following topics:

- Network performance guidance when streaming Provisioning Services disk images to XenServer VMs
- Information about disabling the Spanning Tree protocol

This chapter assumes that you are familiar with how Provisioning Services basic works and its components. For an introduction to Provisioning Services, see the *Provisioning Services Installation and Configuration Guide* for your Provisioning Services version.

Virtualizing the Provisioning Services Server

Before deciding whether or not to host the Provisioning Services server on a VM, it is important to understand the bottlenecks and constraints of Provisioning Services servers, including how Provisioning Services servers use system cache. It is also important to understand that one of the primary resource bottlenecks of Provisioning Services servers is network I/O.

Recommendations

1. Citrix suggests that you host Provisioning Services servers on XenServer VMs that use Single Root I/O Virtualization (SR-IOV). SR-IOV lets a single PCIe device to appear as multiple separate physical PCIe devices. To use SR IOV, the NIC on the host must support SR-IOV. When configured, each VM behaves as though it is using the NIC directly, which reduces processing overhead.



Without configuring SR-IOV, it may be difficult to achieve throughput for Provisioning Services traffic over 2 gigabit speeds.

A Provisioning Services host could under-perform and result in a degraded user experience when there are high numbers of simultaneous VM startups.

To ensure solid network performance, it is critical to carefully consider the write cache configuration. In user workload scenario there are too many users working simultaneously, writing to and reading from the Provisioning Services write cache.

Configuring SR-IOV can mitigate this issue. For configuration information, see CTX126624—[XenServer SR-IOV Support for Provisioning Services Virtual Machines](#).

Note: SR-IOV is only supported on XenServer VMs that are running Windows Server 2008 as the guest operating system.

2. Ensure you are running at least XenServer 5.6 Feature Pack 1. This is to give more the Control Domain more CPU resources. As of XenServer 5.6 Feature Pack 1, the Control Domain can exploit multi-cores and is capable of using up to four cores.
3. Consider increasing the amount of memory allocated to the Control Domains on the hosts running the Provisioning Services VMs. You can increase the memory allocated to the Control Domain from the default allocation of 752MB by following CTX126531—[Configuring Dom0 Memory in XenServer 5.6](#).
4. If you want to virtualize Provisioning Services servers, consider reducing the target device to Provisioning Services server ratio so that the target devices do not exceed 300 targets. While virtualized Provisioning Services servers can support more hosts than this, reducing this ratio can improve performance.

See the Citrix Consulting white paper, CTX128645—[Design Considerations for Virtualizing Citrix Provisioning Services](#).

Note: It is not possible to perform live migration of Provisioning Services server VMs with SR-IOV configured. This is because the SR-IOV NIC's virtual function (VF) is directly tied to a specific virtual machine. (In XenServer, VFs are configured through the XenServer CLI whereas the VM template contains the configurations for the virtual network interface functions and bridges in the Control Domain.) For High Availability design, you can use the Provisioning Services High Availability functionality instead of assuming live migration.

Additional Provisioning Services Networking Considerations

This section discusses some high-level IP addressing requirements and information about isolating the streaming service. There are additional considerations concerning the TCP Large Send Offload option. For more information about all of these topics, see CTX117374—[Best Practices for Configuring Provisioning Server on a Network](#).



IP Addressing Requirements

The Provisioning Server and Provisioning Services Target devices require at least two NICs with IP addresses on different networks:

- One network provides inter communication between devices specifically for the streaming I/O traffic.
- The other network provides access to network resources, the internet, and so on.

Isolating the Provisioning Services Streaming Service

Provisioning Services segment the stream traffic whenever applicable for several reasons: performance, provisioning growth and troubleshooting are a few.

When there are bottlenecks in network capacity or along the backplane of switches, UDP traffic tends to be the first to get dropped or discarded. As a result, Citrix encourages segmentation. When segmented, the streaming service does not have to compete for bandwidth, which ensures performance for the provisioned infrastructure. Segmentation can virtually eliminate retries and maximize Provisioning Services Target/Server performance.

Disabling the Spanning Tree Protocol

When deploying XenServer with Provisioning Services, as a best practice, disable the Spanning Tree Protocol on the switch (that is, set FastPort=On).

Citrix makes this recommendation because PXE takes initialize quickly and, consequently, means that it is best if the switch initializes as rapidly as possible to prevent the PXE environment from timing out. For similar reasons, Citrix recommends disabling the Spanning Tree protocol, since it is slow to initialize.

It is not necessary to hard code the default auto negotiation setting unless you notice long booting times and PXE timeouts.

Best Practice Configurations

Consider the following best practice guidelines when configuring XenServer –Provisioning Services deployments:

1. Follow the guidelines in CTX117374—[Best Practices for Configuring Provisioning Server on a Network](#).
2. Review the information about calculating IOPS for Provisioning Services servers in CTX125126—[Advanced Memory and Storage Considerations for Provisioning Services](#). **Note:** The guidance about CIFS no longer applies if you are running Windows Server 2008 R2.



3. If you want to deploy Provisioning Services on blades, review the Citrix Community Blog Post, "[Optimizing PVS.](#)" This article describes how to confine most Provisioning Services network traffic in the blade chassis so as to enable high-performance network connections between the blades.

Chapter 8: Verifying Your XenServer Networking Configuration

This chapter provides some basic steps for verifying that you configured and physically connected XenServer networking correctly. This chapter includes:

- Steps for checking your configuration
- An overview of steps you can take to resolve issues

Verifying XenServer Networking on a Pool

This chapter provides you with ways to verify your networking configuration before you create VMs and deploy the pool. Checking to see if your networking is configured properly before you proceed with further configuration can reduce troubleshooting and reconfiguration later.

Create a resource pool and run the tests in this chapter from any host in the pool. The tests in the first topic, “Verifying your Physical Configuration,” verify your physical configuration, including if your NICs are fully functioning. The tests in the other major topic, “Verifying your XenServer Networking Settings and Configuration,” verify whether your XenServer networking settings are configured correctly.

Verifying your Physical Configuration

This section provides procedures to help you verify NIC connectivity and NIC speed. For testing purposes only, create the following:

- A Windows VM and a Linux VM with XenServer Tools installed on them.
- An external network.



- An external bonded network.
- A XenServer single-server private network. Add two or more VMs (with XenServer Tools installed on them) to that network. You will need to manually configure IP addresses for these VMs unless a machine running DHCP was added to the private internal network.

Issues with NIC connectivity and NIC speed can indicate hardware that is not compatible with the [XenServer Hardware Compatibility List](#). They can also indicate issues with the NIC driver.

Verifying NIC Connectivity

1. Using Windows Remote Desktop Connection or a similar tool, verify that it is possible to connect to the Windows VM over both an external network and an external bonded network.
2. Using Iperf, verify connectivity between the VMs on the XenServer single-server private network.

Verifying NIC Speed

Testing to see if the NIC can send and receive traffic at the approximate speed for which it was rated can reveal issues, such as, for example, problems with the NIC drivers.

Before beginning the tests, download and install Iperf on a Windows VM, a Linux VM, and a separate physical server that is not running XenServer (a *control host*).

Run both the TCP bi-directional test and the UDP tests between each VM and the control host. For example, run the first test between (1) the Linux VM and the control host and (2) the Windows VM and the control host.

TCP Bi-Directional Test

1. On the Windows VM, run the following command with the duration set to 300 seconds:

```
iperf.exe -c 192.168.1.1 -d -t 300
```
2. While running the command on the VM, run the following command on the control host:

```
iperf.exe -s
```
3. Repeat steps 1 and 2 using the Linux VM instead of the Windows VM.
4. Repeat steps 1 to 3 using a NIC bond.



UDP Test

1. On the Linux VM, run the following command with the duration set to 300 seconds. In this test, the VM functions as the Iperf client.

```
iperf.exe -c 192.168.1.1 -u -t 300 -b <bandwidth>
```

Replace <bandwidth> with “100M” for a 1 gigabit NIC. For a 10 gigabit NIC, replace <bandwidth> with “1000M”.

2. While running the command on the VM, run the following command on the control host:

```
iperf.exe -s -u
```

3. Repeat steps 1 and 2 using the Windows VM instead of the Linux VM.
4. Repeat steps 1 to 3 except using each VM as the Iperf server and the control host as the Iperf client.
5. Repeat steps 1 to 4 using a NIC bond.

Verifying your XenServer Networking Settings and Configuration

After setting up your XenServer networking configuration, verify that you did it correctly by checking the following:

1. Management and Secondary Interfaces (Managed Interfaces).

Configure networking after forming the pool, as described in “Sequence of Networking Configuration Tasks” on page 18. After creating the initial networking configuration or adding hosts to a pool, it is possible for the managed interfaces on the other hosts in the pool to be incompletely configured (for example, if the pool master does not successfully propagate its settings to the member or joining servers).

1. Verify that each management and secondary interface uses the same (corresponding) NIC on each XenServer host.
 2. Verify that you set a unique IP address on each XenServer managed interface or bonded set of managed interfaces.
 3. Verify a) and b) on any secondary interfaces you configured (for example, for storage)
2. Verify all the NICs you bonded were created correctly. To do so, check the XenCenter logs for failures or any bonds that unfinished. A typical sign that bonding did not work is that the bond takes excessively long to finish or it seems to “hang.”

3. Verify that the networking configuration and physical cabling on each XenServer host matches. For example, is NIC 3 on Host 4 configured with the same networks and cabling as NIC 3 on Host 5.

Tip: To make an LED light for a NIC on a XenServer host blink, run the **ethtool** command with the **-p** option on the host. **ethtool** is a Linux command that displays or changes NIC settings. For example, to use the **ethtool** command to make the lights on the first NIC on the host, run the following on the host: **ethtool -p eth0 10**. In this example, **eth0** is the NIC and **10** is the number of seconds you want the NIC to blink.

4. In XenCenter, select each host individually, and verify the information in the **NICs** and **Network** tabs is fully present.

In the **NICs** tab, there should be a row in the NICs list for each NIC. Verify the columns in the list have correct values, especially the speed and MAC address columns.

In the **Networks** tabs, verify that all the NICs listed in the **NICs** tab are present and they have the correct network listed beside them.

5. Verify that the physical networking configuration matches for each XenServer host (that is, check to make sure that NIC 3 on Host 4 is cabled to have the same connectivity as NIC 3 on Host 5 and so on).
6. If you are using a vSwitch as the networking bridge, verify that it is enabled on all hosts in the pool.
7. Install two VMs on different hosts, if necessary, and put them on the same network to see if they can “ping” each other.

Resolving Issues

If, after checking these items, VMs or hosts cannot connect to each other as expected, then:

1. Verify that your physical configuration (cables and switches) is correct.

Important: The #1 networking issue the XenServer technical support team sees is not actually a *XenServer* issue. Most calls are from customers with incorrectly configured physical networks (cabling errors) or incorrectly configured switches. Before calling Citrix Technical Support, verify your cabling configuration.

1. Review the section, “Recovering from a bad network configuration” in the *XenServer Administrator’s Guide*.

Important: After the pool is in production, always use extreme care when changing the physical networking hardware in a host. Before changing, adding, or replacing any NICs on the host,



always take the VMs offline and put the host into maintenance mode so that the host is no longer connected to any storage repositories.



Revision History

Revision	Comments	Date
1.0	Initial Release	April 28, 2011
2.0	Updated for XenServer 6.0 release.	September 23, 2011
2.1	Made minor correction to syntax of xe-switch-network-backend command. (Hyphen missing between xe and switch.)	October 20, 2011
3.0	Updated for the XenServer 6.1 release. Updates include information about LACP bonding and networking for Storage XenMotion. This update also include minor terminology changes.	September 2012



About Citrix

Citrix Systems, Inc. (NASDAQ:CTXS) is the leading provider of virtualization, networking and software as a service technologies for more than 230,000 organizations worldwide. Its Citrix Delivery Center, Citrix Cloud Center (C3) and Citrix Online Services product families radically simplify computing for millions of users, delivering applications as an on-demand service to any user, in any location on any device. Citrix customers include the world's largest Internet companies, 99 percent of Fortune Global 500 enterprises, and hundreds of thousands of small businesses and prosumers worldwide. Citrix partners with over 10,000 companies worldwide in more than 100 countries. Founded in 1989, annual revenue in 2010 was \$1.87 billion.

©2012 Citrix Systems, Inc. All rights reserved. Citrix®, Access Gateway™, Branch Repeater™, Citrix Repeater™, HDX™, XenServer™, XenCenter™, XenApp™, XenDesktop™ and Citrix Delivery Center™ are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other trademarks and registered trademarks are property of their respective owners.