40555A Networking Fundamentals

Module 4: Extending your network

Contents

Learning objectives based on Microsoft Technology Associate (MTA) exam objectives

#	Lesson title	Learning objectives	Exam objectives mapped
1	Network switches	• Describe a switch. • Describe virtual local area networks (VLANs). • Explain virtual switches.	1.2.4 VLANs 2.1.1 Transmission speed 2.1.2 Number and type of ports 2.1.3 Number of uplinks 2.1.4 Speed of uplinks 2.1.5 Managed or unmanaged switches 2.1.6 VLAN capabilities 2.1.7 Layer 2 and Layer 3 switches and security options 2.1.8 Hardware redundancy 2.1.9 Support 2.1.10 Backplane speed 2.1.11 Switching types and MAC table 2.1.12 Understand capabilities of hubs versus switches 2.1.13 Virtual switches

#	Lesson title	Learning objectives	Exam objectives mapped
2	Routing	• Describe routing. • Compare static and dynamic routing. • Explain routing tables.	2.2.1 Transmission speed considerations 2.2.2 Directly connected routes 2.2.3 Static routing 2.2.4 Dynamic routing (routing protocols) 2.2.6 Default routes 2.2.7 Routing table and how it selects best route(s) 2.2.8 Routing table memory 2.2.12 Quality of Service (QoS)
3	Routing protocols	• Describe Routing Information Protocol (RIP). • Describe open shortest path first (OSPF).	2.2.5 RIP vs. OSPF
4	Routing in Windows Server	• List available routing options in Windows Server. • Explain how to implement routing in Windows Server. • Describe network address translation (NAT). • Explain how to implement NAT.	2.2.9 Network Address Translation (NAT) 2.2.10 Software routing in Windows Server 2.2.11 Installing and configuring routing 3.5.2 Network Address Translation (NAT)

Module overview

In this module, we'll examine the technologies that enable you to extend your network. Precisely what technology you use will depend on how far and in what way you want to extend it. For example, some network topologies have imposed limits on the number of devices that can be connected to a single wiring concentrator. Others impose constraints on the distances over which you can reliably transmit data between networked nodes. Essentially, you can extend your network in two ways: by using network switches, or by using routers and selecting an appropriate routing protocol.

Objectives

After completing this module, you will be able to:

- Describe network switches.
- Explain routing.
- Describe common routing protocols.
- Explain how to implement routing in Windows Server.

Lesson 1: Network switches

In this lesson, you'll learn about devices that enable you to connect host computers together. At the basic end of the spectrum, you can use a simple hub, which in essence is a wiring concentrator. At the other end of the spectrum, you can use a switch that can implement bridging and routing functionality to control the flow of traffic in your network infrastructure.

Objectives

After you complete this lesson, you will be able to:

- Describe a switch.
- Describe VLANs.
- Explain virtual switches.

What is a switch?

A *switch* is a device with physical ports that enable wired connections from your networked devices and hosts, as Figure 1 depicts. In some respects, a switch seems like a hub.

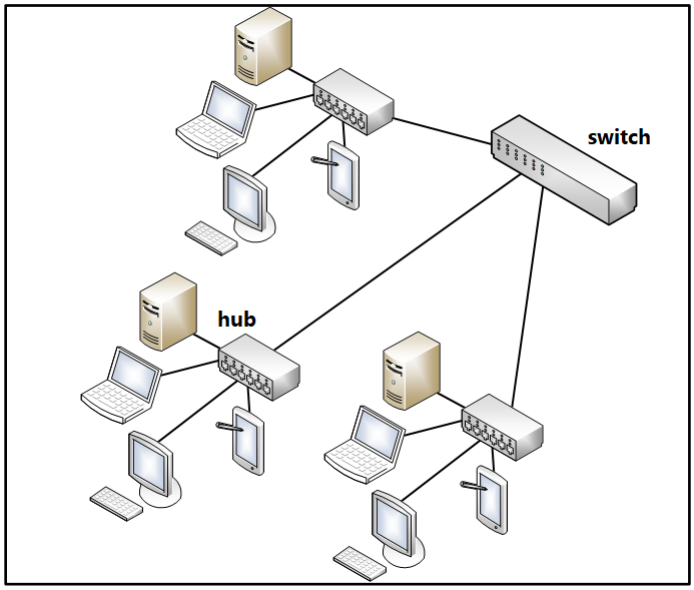

Figure 1. Local network that's wired by using a switch

Note

In Module 2, "Local area networks and wide area networks," we talked about local area networks (LANs). We discussed Ethernet, and you learned that Ethernet was a logical bus topology that's typically star wired. Early Ethernet networks were wired by using a single length of coaxial cable, and each device was connected to it by using a drop cable.

When you use a hub, you essentially collapse the long coaxial segment into the hub, and you extend the drop cables from the hub in a star configuration to each device.

But a hub is a basic wiring concentrator. When you use a hub, you're simply connecting devices together in a star-wired arrangement. There's no management of network traffic; if you use hubs, your network is effectively one big segment, one single Internet Protocol (IP) subnet.

You'll remember that Ethernet, the most widely used network topology, is a contention system in which devices compete for available bandwidth. Thus, the more devices you connect, the less effective the available bandwidth is, as Figure 2 depicts. This is because the devices spend most, if not all, of their time managing collisions and retransmissions of data.

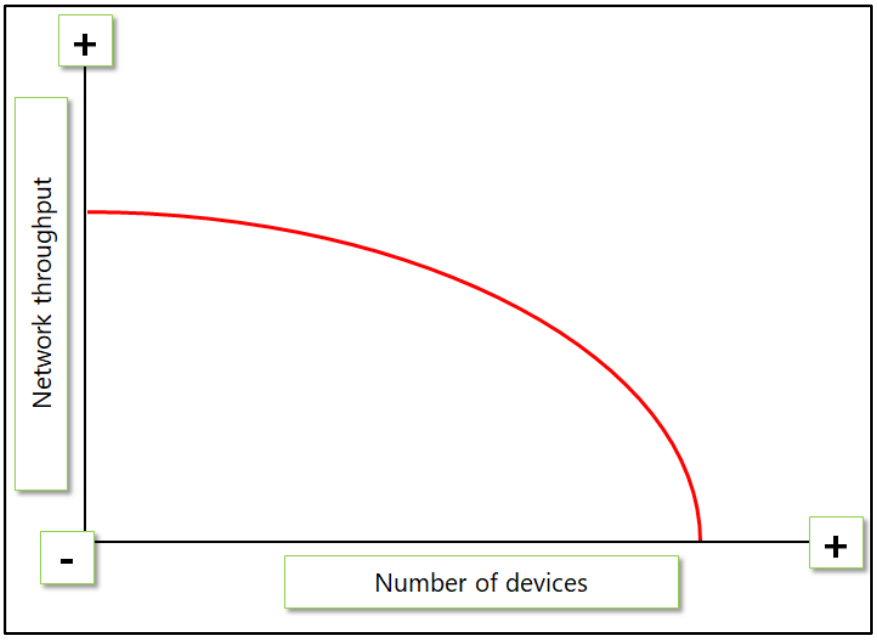

Figure 2. Network throughput performance by number of devices

However, if you reduce the number of devices, effective bandwidth use (network throughput) increases. One way to manage network bandwidth is to use switches instead of simple hubs.

Note

Although switches are most commonly available for Ethernet networks, they're also available for other technologies, including asynchronous transfer mode (ATM) and Fibre Channel.

The key difference between a hub and switch is that a hub repeats electrical signals onto each port without concerning itself with the nature of the traffic and its destination. A switch, however, decides the destination of the traffic, and it doesn't always forward traffic to all connected ports. Rather, it forwards traffic only to ports that have connected devices that require the data.

Switches are available to operate at a number of different levels. The level of operation determines the functionality of the switch. Additionally, some switches offer managed functionality. The following sections discuss these options.

Switching types

Switches are available to operate at different levels of the Open Systems Interconnection (OSI) model, as Figure 3 depicts:

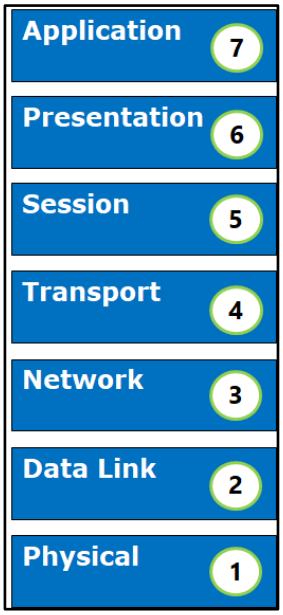

Figure 3. The OSI model

Switches are, therefore, often described as layer 2, layer 3, or layer 4 switches.

Note

Layer 7 switches work at the application layer and distribute traffic based on selected Uniform Resource Locators (URLs).

Layer 2 and layer 3 switches

The most common types of switches are layer 2 and layer 3. Let's examine these in more detail.

Layer 2 switches

A switch that operates at layer 2 operates at the media access control (MAC) layer (layer 2) by connecting networks together and passing frames between them.

Layer 2 switches operate in *promiscuous mode* on attached networks. This means they receive all network frames on all connected interfaces and then forward all those frames to all other interfaces. Layer 2 switches aren't specifically addressed by hosts, which means that a layer 2 switch isn't visible to a host.

Layer 2 switches have the following key features:

- They operate at the data-link layer (MAC layer) and are, therefore, protocol independent.
- They're transparent to communicating hosts. In other words, they don't realize that their frames are traveling through a layer 2 switch to get to the destination host.
- They always pass broadcast and multicast traffic.

In essence, a layer 2 switch forwards all broadcast and multicast traffic. However, it only forwards unicast traffic to specific devices based on device MAC address.

To some extent, this means that a layer 2 switch can help localize network traffic. Only required traffic passes between ports. To manage network traffic properly, however, you must segment a network into logical IP subnets. A layer 2 switch can't do that.

Note

You might be wondering how a layer 2 switch knows which devices are accessible on which of its ports. The answer is that the switch learns by examining the source MAC address of all frames it receives. If it receives a frame on port 23 that has a source MAC address of 02608C123ABC, it knows that the device with the MAC address of 02608C123ABC is accessible on port 23.

Layer 3 switches

A layer 3 switch behaves like a router. A router—and therefore a layer 3 switch—operates at the network layer of the OSI model (the internet layer in a TCP/IP network). Unlike bridges and layer 2 switches, hosts specifically address routers. When a host wants to communicate with a device in another network (subnet), it must have a route to that subnet. Failing that, it sends its packets to its configured router, which is sometimes called a *default gateway*.

The router is responsible for onward forwarding. Unlike bridges, routers only propagate network packets that are specifically addressed to a remote subnet; this enables routers to help manage network traffic by localizing subnet traffic.

While a layer 3 switch is like a router, it's not identical. For most routing needs in a LAN environment, however, a layer 3 switch provides the required network traffic management. Most layer 3 switches support the notion of routing between VLANs. The next topic will cover VLANs, and we'll also discuss more about routers and routing in the next lesson.

Managed vs. unmanaged switches

Unmanaged switches have no management interfaces or configuration options. They're simple devices that you plug into and turn on. Because you can't manage and configure them centrally and remotely, they're not ideal for large workplace networks. However, they're relatively inexpensive and are often deployed in small offices and home office environments.

Managed switches enable a network administrator to remotely interface with the switch and reconfigure it. Typically, switches support a command-line interface that's accessible by a local connection and cross-network interfaces through Telnet or via a Simple Network Management Protocol (SNMP) agent.

Note

SNMP is a TCP/IP standard for network management. Devices with an agent are manageable from a management console or tool that also adheres to the SNMP standard. Standards enables an organization to use tools and devices from different vendors and still be able to manage them without using proprietary apps and interfaces.

Typical management capabilities might include:

- Configuring VLAN settings.
- For layer 2, configuring MAC address tables and filtering options.
- Managing bandwidth settings across available ports.
- Enabling or disabling ports.

Characteristics of switches

When deciding whether to use switches and choosing what type of switch to deploy, you must consider many factors, including:

- Transmission speed. A critical factor in determining which switch to use is its speed. Broadly, there are two transmission speeds: Fast Ethernet and Gigabit Ethernet. The former supports speeds of up to 100 megabits per second (Mbps), the latter supports speeds of 1,000 Mbps, or 1 gigabit per second (Gbps). The speed that you choose is determined by the bandwidth you anticipate requiring for a particular network. Fast Ethernet is considered suitable for most LAN situations and is used for connecting users' devices and peripherals, such as in a single department. Gigabit Ethernet supports ports to which servers are connected or LANs that require additional throughput.
- Backplane speed. The *backplane speed* of a switch is the bandwidth of the network within the switch that connects the switch's various ports—in other words, the maximum volume of data the switch can handle.

- Number and type of ports. Larger networks require switches that have more connections, or *ports*. The different types of ports on most switches include:
    - RJ45. These ports connect LANs by using copper over unshielded twisted pair (UTP) media. You must have sufficient RJ45 ports to accommodate the number of routers, servers, and downstream switches and hubs.

Note

As discussed in Module 3, "Media types," there are a number of copper UTP cabling standards. You must ensure that your switches and wiring adhere to the exact same standard, such as Cat 5, Cat 5e, and Cat 6.

    - SFP. Small form-factor pluggable (SFP) ports support connections from transceivers that connect to Fibre Channel or Gigabit Ethernet. Most high-end switches come with at least some SFP ports.
    - SFP+. An improved version of SFP, these ports support 10-Gbps speeds and can be used to connect to 10-GB switches.
- Number of uplinks. Uplinks connect a switch to the network infrastructure. For example, if you use a Fast Ethernet switch to connect the computer devices in a department (laptops, printers, servers, and so on), then you use the uplink port to connect the departmental switch to the backbone.
- MAC tables. Also known as *content addressable memory* tables, switches use these to identify hosts' locations on connected ports.
- Hardware redundancy. If you deploy a switch and connect all your devices to it, then the switch becomes a single point of failure. A maxim of any system design is to always avoid single points of failure. One possible solution is to deploy switches that support redundancy systems. Link aggregation is another possible solution. The Institute of Electrical and Electronic Engineers (IEEE) 802.3ad Link Aggregation Control Protocol (LACP) enables parallel switch-to-switch connections. If redundancy is critical in an organization, using LACP-compliant devices might mitigate single points of failure.

Note

Redundancy in layer 2 switches can be problematic. Remember that the layer 2 switch uses MAC tables to learn where devices are. The source MAC address 02608C123ABC in an Ethernet frame that originates on port 23 indicates to the switch that the host with the MAC address of 02608C123ABC resides on port 23. However, if you build redundancy into your design, you can create loops. This might mean that a frame originating from the host with the MAC address 02608C123ABC appears on port 12, having been looped around a failure. This can confuse the switch. To avoid this issue, consider implementing switches that support the spanning tree algorithm when planning redundancy in layer 2 switches.

What are VLANs?

A virtual LAN (*VLAN*) is pretty much what the name says—it's a LAN, virtually. In other words, by using a technique known as *frame tagging*, switches can create the appearance of multiple LANs interconnected by routers, to manage network traffic, as Figure 4 depicts:

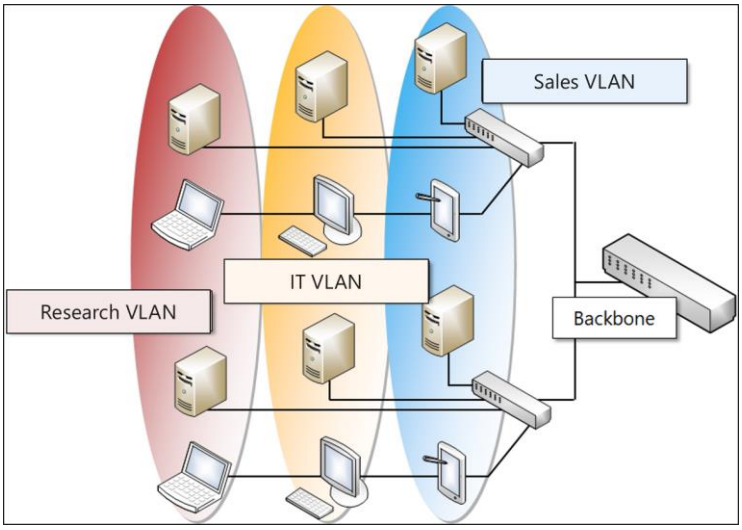

Figure 4. Implementing VLANs

The major advantage of VLANs is that network administrators can create logical networks of various hosts and devices that aren't even connected to the same switch.

For example, to create the Sales VLAN in Figure 4, without VLANs, the administrator would need to create a backbone network and connect the various switches necessary for the Sales users to the backbone. Given that these Sales users are distributed throughout the building on various network segments, that would be time-consuming and expensive. Instead, the administrator identifies to which ports the sales switches are connected, and using software, reconfigures those switch ports as part of a virtual network. There's no re-cabling, and it's a comparatively simple and quick change to make.

Note

Less expensive switches that support VLANs might do so at the port level only. That is, you can mark a port as being a VLAN. You can't mark only some of the devices on that port as being part of a VLAN. To mark specific hosts only, you must implement a switch that supports VLAN frame tagging. The IEEE 802.1Q standard defines this.

To be precise, a VLAN represents a *broadcast domain*. That is, the collection of devices that would receive an Ethernet broadcast frame. By defining VLANs, you effectively create multiple routed networks; in fact, a VLAN maps directly to an IP subnet.

As Figure 4 illustrated, VLANs can cross physical network component boundaries to create logical networks that span (in this case) departments that are distributed throughout a building.

A significant benefit to using VLANs for a network administrator is that they can be relatively easily reconfigured, whereas physical subnets might well require cabling and infrastructure changes.

Virtual switches

Virtual switches are found in computer operating systems that support virtualization. With virtualization, a computer can share its physical hardware with one or more isolated operating systems that are running in virtualized environments on virtual machines (VMs).

VMs share the physical resources of a physical computer, and they represent those virtualized resources as usable components to the operating systems that are running on the VMs.

A virtual switch provides the virtual equivalent of a physical switch. It enables you to define which VMs are connected to the host, to each other, or which are isolated. You can also define advanced characteristics for your virtual switches to control bandwidth and other settings.

Windows 10 supports a feature called *Client Hyper-V*, which is a visualization platform. Client Hyper-V supports three types of virtual switches, which Figure 5 illustrates:

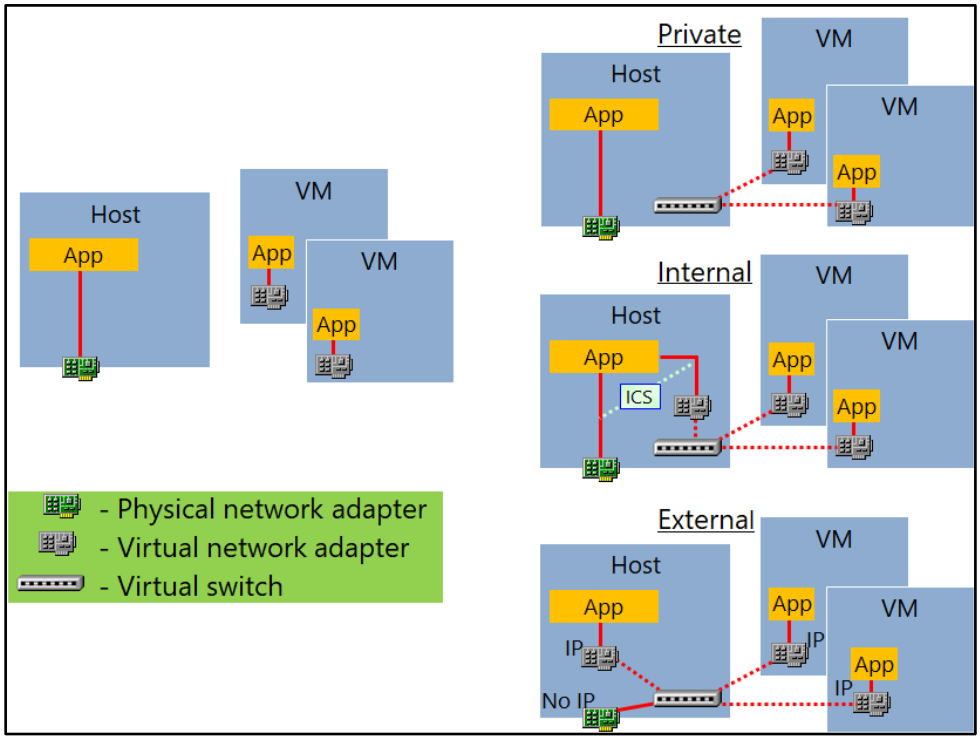

Figure 5. Implementing virtual switches

The three types of virtual switches that Client Hyper-V supports include:

- Private network. A virtual switch that you connect to a private network provides connectivity only between VMs that are running on the same Client Hyper-V computer, and that connect to the same virtual switch. VMs can't communicate with VMs that are connected to a different virtual switch, a physical Windows 10 computer, or to an external physical network. You can use a private switch if you need to isolate VMs for security reasons or if you use them for testing and you don't want them to inadvertently access the organizational network.

- Internal network. A virtual switch that you connect to an internal network that provides connectivity between VMs that are running on the same Client Hyper-V computer and on the Client Hyper-V computer itself. VMs that are connected to an internal switch can't communicate with any physical network unless the physical Windows 10 computer provides the Internet Connection Sharing (ICS) functionality or if you create a NAT object.

  You use an internal virtual switch when you want VMs to have network connectivity with a physical Windows 10 computer, but you don't want them to access external resources.

- External network. A virtual switch is connected to a physical network adapter or to a wireless adapter on the physical Windows 10 computer. This switch enables VM connectivity to a physical network. You use an external switch to provide VMs access to external resources or to the internet.

Lesson 2: Routing

Routers are widely implemented within organizations to interconnect subnets (or VLANs) in large networks and to interconnect LANs to create wide area networks (WANs). In this lesson, you'll learn about routing fundamentals.

Objectives

After you complete this lesson, you will be able to:

- Describe routing.
- Compare static and dynamic routing.
- Explain routing tables.

What is routing?

Routing is the process a router undertakes on behalf of a host (or another router) to ensure the onward delivery of a packet that's destined for a remote host—that is, a host on a remote network.

Routers are internetwork devices that are connected to multiple network interfaces. When a router receives a network packet, it determines the best way to forward the packet to its destination. Routers are internetwork store and forward devices, which differs from bridges in several important ways, including that:

- They operate at the network layer.
- They're protocol dependent.
- They're addressed by the hosts on the network.
- They can optimize network availability by managing multiple routes effectively.
- They can connect LANs together that differ in their topology, such as token ring to Ethernet.

- They provide more reliable, available networks (no need to worry about the spanning tree algorithm).

Note

Windows Server includes Routing and Remote Access Service support, which enables routing over IP networks either by connecting LANs or by connecting LANs to WANs without needing to purchase a dedicated router.

Local vs. remote

One of the critical things a host does is to determine whether a packet needs to be routed or whether the packet can be delivered locally. Let's examine this process, which Figure 6 illustrates:

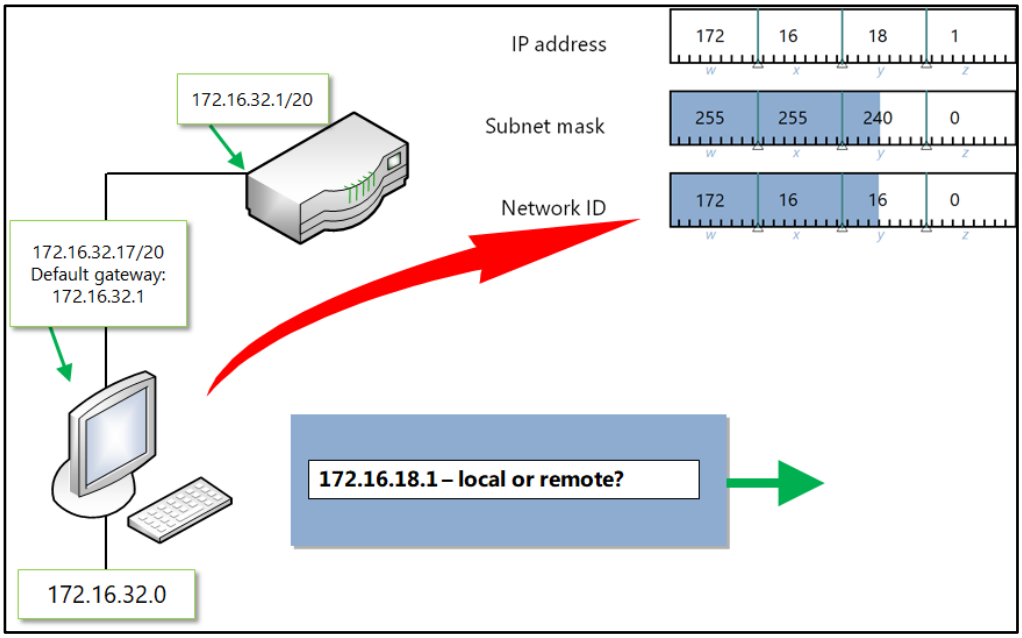

Figure 6. Determining whether to route

Local delivery

Network communication typically begins with a connection request to another host by its computer name. The sending host determines whether the target host is local (same subnet) or remote (different subnet). If local, it broadcasts an Address Resolution Protocol (ARP) request for the target host's MAC address, and based on the response, it sends the required Ethernet frames to the target host.

Remote delivery

In remote delivery, after the sending host determines that the target host is in a different subnet, it broadcasts an ARP request for its gateway's MAC address, and then it sends the required Ethernet frames to the gateway. The gateway then routes the packets contained in the frames to the appropriate subnet.

The following steps are an overview of the process that Figure 6 depicts:

1. A host sends a request to connect to Server1. The name Server1 is resolved to an Internet Protocol version 4 (IPv4) address. (A later module in this course will discuss name resolution.) Let's suppose the IP address is 172.16.18.1.

2. After the sender knows the recipient's IPv4 address, it uses its subnet mask to determine whether the IPv4 address is remote or on the local subnet. The subnet mask in Figure 6 uses 20 bits, which is 255.255.240.0. This means that the host, Server1, is in a different subnet (subnet 172.16.16.0/20). The local host is in subnet 172.16.32.0/20.

3. Because Server1 is in a different subnet and is therefore remote, the local host sends an ARP request to its configured default gateway (172.16.32.1). The gateway responds to the ARP request with its MAC address.

4. The sending host forwards the packets for Server1 to its gateway. It does this by creating Ethernet frames that contain the fragmented packets, and it then merges these frames onto the wire for delivery to the MAC address of the router.

5. The router receives the frames, strips off the headers, and examines the encapsulated packets. It then routes those packets to the correct subnet.

Routing protocols

As mentioned earlier, routers are protocol dependent. However, routers support different routing protocols depending on the purpose of the router. For example, routers that connect a small number of LANs use a certain type of routing protocol. But routers that connect large organizations to the internet use a different type of routing protocol.

Note

Routing protocols enable a router to learn about the network (or networks) to which it's connected and to identify changes in those networks. This is analogous to understanding routes from home to your workplace. You might know of three or four different routes. But let's say that one day, you learn about road resurfacing work, and you decide to use an alternate route. That's what routing protocols are all about—learning and maintaining routes.

A routing protocol has three basic jobs to perform:

- Discovery. A router must be able to discover, or *learn about*, new routes.
- Management. A router must know about all the various destinations for packets and have a way to identify the characteristics of the path to each.
- Path. A router must be able to decide which route to take.

The following table describes the different routing protocols and their purposes.

Protocol	Explanation and purpose
Routing Information Protocol (RIP)	An early routing protocol designed predominantly for smaller networks, it's somewhat limited for larger networks. We'll examine RIP in more detail in the next lesson.
Open shortest path first (OSPF)	Designed to overcome RIP limitations, it's widely supported by different vendors. We'll

Protocol	Explanation and purpose
	examine OSPF in more detail in the next lesson.
Border Gateway Protocol (BGP) and External Gateway Protocol (EGP)	Used by internet providers and larger organizations to identify and distribute changes to routing tables within network infrastructure, it's quite complex. BGP is the internet version of EGP.
Interior Gateway Routing Protocol (IGRP) and Extended IGRP (EIGRP)	IGRP was developed by Cisco to provide another alternative to RIP. Replaced by EIGRP, the new standard supports Classless Interdomain Routing (CIDR).

Static and dynamic routing

You can configure a router with the necessary information to route packets in several ways:

- Configure a router with a default gateway address, and when it doesn't have a route, it will forward the packet to the designated host.
- Configure a routing table or a list of entries for routes to other networks. This is known as *static routing*.
- Install a routing protocol such as RIP or OSPF. These protocols learn about routes from other routers. This is known as *dynamic routing*.

Choosing between static and dynamic routing depends on your network and routing needs. The differences between static and dynamic routing include:

- Static routing is a function of the IP protocol.
- Routing tables must be built and maintained for static routing to function.
- Static routers must be reconfigured in the event of any internetwork reconfiguration.
- Static and dynamic routers don't communicate with each other.

- Routers don't typically pass either broadcasts or multicasts.
- Windows Server computers that are configured with multiple network adapters can function as multihomed computers and will support both static and dynamic routing.
- To route IP packets, a static router must have either:
  - An entry for each network on the internet in their routing table.
  - The default gateway address must be configured with another router's local interface address.

Note

Using the default gateway method is only effective for two routers.

Routing tables

As we have already learned, routers must know about routes to other networks. Whether you manually enter routes (static) or use a routing protocol (dynamic), the routes are stored in a routing table, as Figure 7 depicts:

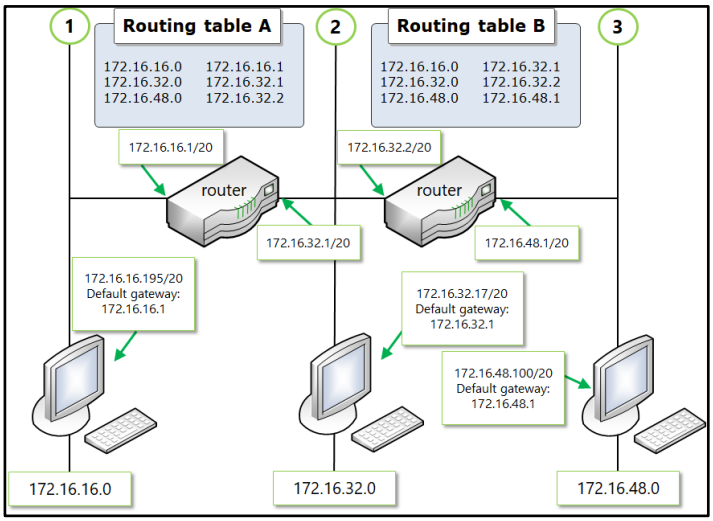

Figure 7. Routing tables

In Figure 7, the two routing tables consist of information about the available networks and how to get to those networks. Some routing tables also include information about the cost of getting to a network; this is so that the router can make a sensible decision about which route to take.

Note

Some routers and their routing protocols support updates to routes based on changes in route availability. For example, if a route is up but congested, it might be sensible to use an alternate route.

You can use the **Route.exe** command-line tool on a Windows 10 computer to observe and maintain routing information. The following table lists the available options.

Command	Explanation
Route add [*network*] mask [*netmask*] [*gateway*]	Adds a route
Route –p add	Adds a persistent route (added to the registry)
Route delete [*network*][*gateway*]	Deletes a route
Route print	Displays the routes
Route change [*network*][*gateway*]	Changes an existing route

Note

Unless a Windows computer is installed with multiple network adapters and is configured as a software router (as discussed in the last lesson), you generally shouldn't need to maintain routes on a host. This is because in the absence of a configured route, a computer uses its configured default gateway as the preferred route.

Default routes

All hosts maintain routing tables—even a Windows 10 computer has a routing table. The entries in a host's routing table are often known as *default routes*. Figure 8 depicts the default routes on a Windows 10 computer:

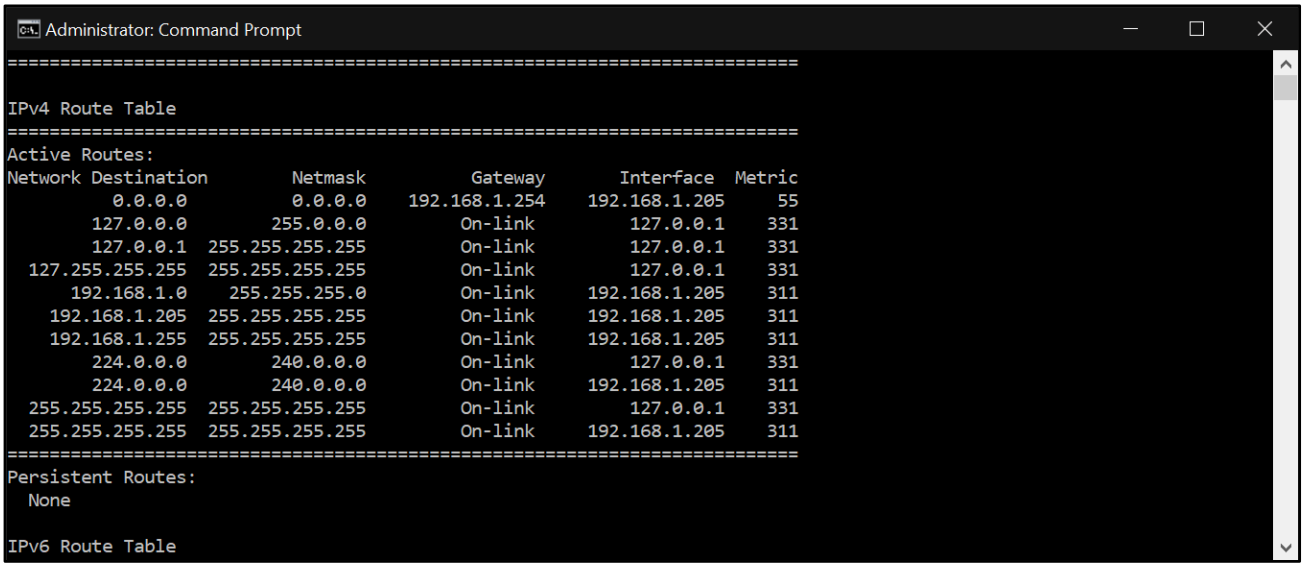

Figure 8. Adding routes at the command line

The following routes are identified for a host with the IPv4 address of 192.168.1.205/24 and a default gateway of 192.168.1.254:

- 0.0.0.0. This is the default route, which is accessible from the local interface, and it uses the default gateway as a gateway address.
- 127.0.0.0. This is the loopback subnet. Loopback is used for testing purposes and can't be assigned to a host.
- 127.0.0.1. This is the IP address that's allocated to the loopback adapter.
- 127.255.255.255. This is the broadcast address for the local loopback subnet.
- 192.168.1.0. This is the local subnet address, which will vary depending on your IPv4 configuration.
- 192.168.1.255. This is the local subnet broadcast address, which also varies based on what IPv4 address is assigned to the local subnet.
- 224.0.0.0. This is a multicast address.
- 255.255.255.255. This is a broadcast address.

Lesson 3: Routing protocols

As mentioned earlier, routing protocols are used to configure and maintain routing tables by providing information about available routes between network destinations. In this lesson, we'll compare RIP and OSPF, two common routing protocols.

Objectives

After you complete this lesson, you will be able to:

- Describe RIP.
- Describe OSPF.

RIP

RIP has been around for some time and is ideal for smaller networks. There are two versions, RIP version 1 and RIP version 2. RIP-based routers broadcast (or multicast) routing data to each other at defined intervals. RIP is relatively easy to configure, and it's widely supported.

The following points summarize the protocol's abilities:

- RIP uses User Datagram Protocol (UDP) port 520.
- RIP uses a hop count field, or *metric*, to indicate the distance to the remote network.
- Each hop represents a router through which the packet must pass to reach the remote destination.
- The hop count reaches a maximum after 15 hops.
- Hop numbers of 16 and greater are considered unreachable.

RIP versions

The two versions of RIP are RIP version 1 and RIP version 2, and their key differences are:

- RIP version 1 supports classful routing.
- RIP version 1 uses broadcasts to propagate routing information.
- RIP version 2 supports multicast announcements.
- RIP version 2 supports simple password authentication between routers.
- RIP version 2 supports CIDR and Variable Length Subnet Masks (VLSMs).

You can choose to implement either RIP version 1 or RIP version 2 on Windows Server depending on your requirements. In either case, you install the same RIP protocol, but your configuration options determine the version that you implement.

Problems with RIP

Because RIP uses a distance vector protocol, each router's route table has a complete list of all possible routes to all destination networks. This table could be huge. Because routers communicate by using 512-byte packets, changes to the route table would be propagated by using a vast number of 512-byte chunks.

RIP routers use broadcasts (although RIP version 2 routers can be configured to use multicasts) to propagate their route tables on all attached interfaces every 30 seconds. This can be undesirable on large networks because these broadcasts can consume network bandwidth.

If a router fails, the change in the internetwork doesn't propagate for some minutes because of a problem known as *slow convergence*. This can lead to the loss of packets.

OSPF

OSPF is typically used in very large networks. It's far more complex to configure than RIP, but it's also far more efficient in big internetworks.

The key features of OSPF are that:

- The 16-hop RIP limitation doesn't exist.
- Routing tables update more efficiently. A map called the *link-state database* is created of all the routes. This map then syncs with other routers after changes have occurred in the routing topology.
- The routing environment of an OSPF-based network is more structured, consisting of an autonomous system made up of OSPF areas that all share a common administrative authority.

Lesson 4: Routing in Windows Server

In Windows Server 2016, Routing and Remote Access Service (RRAS) can function as a software-based router and thereby can manage the data that flows between subnets. Its routing capabilities include LAN-to-WAN and NAT.

A router manages outgoing and incoming data packets based on the information in its routing table. It directs the traffic to either the destination or to another router, which then processes the packet and forwards it to the destination network. The routing table contains information about the router´s own network interfaces, destinations, and sources for network traffic.

Objectives

After you complete this lesson, you will be able to:

- List available routing options in Windows Server.
- Explain how to implement routing in Windows Server.
- Describe NAT.
- Explain how to implement NAT.

Routing options in Windows Server

Windows Server 2019 can act as a router or NAT device between two internal networks or between the internet and an internal network, as Figure 9 depicts. Routing works with routing tables and supports routing protocols such as RIP version 2, Internet Group Management Protocol (IGMP), and Dynamic Host Configuration Protocol (DHCP) relay agent.

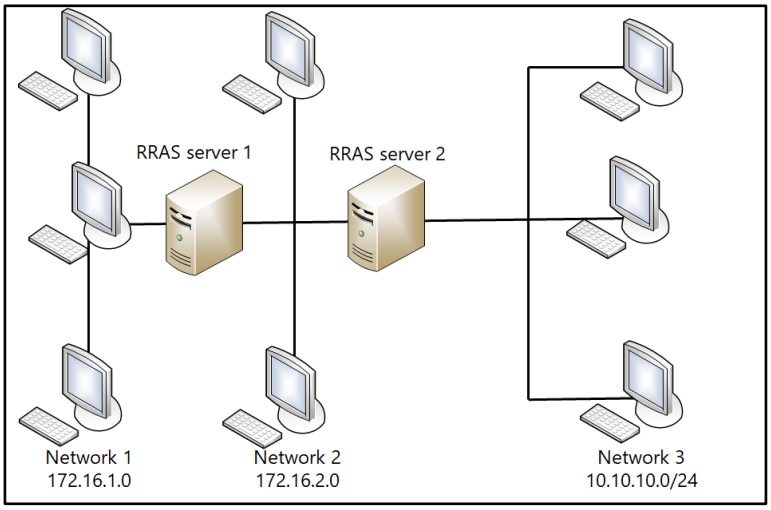

Figure 9. Routing with Windows Server

RRAS supports the following types of routing, as Figure 10 depicts:

- Demand-dial interface VPN routing
- Static routes for both IPv4 and Internet Protocol version 6 (IPv6)
- IGMP for IPv4
- RIP for IPv4
- NAT for IPv4

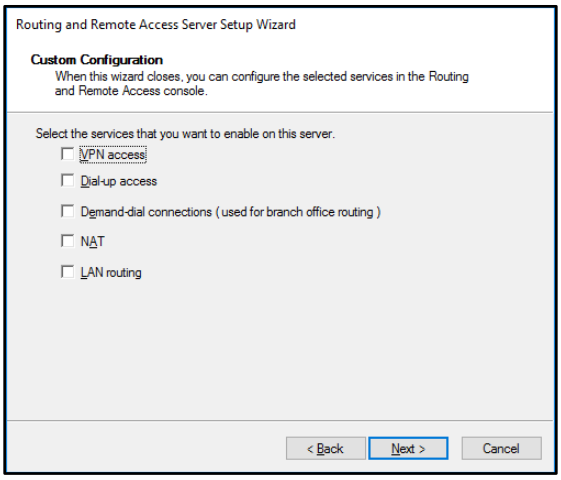

Figure 10. Configuring a software router

Demonstration: Implementing routing in Windows Server

Your instructor will now demonstrate how to deploy and configure a router in Windows Server.

Install the Routing role service

1. Sign in to the appropriate server as **Adatum\Administrator** with the password **Pa55w.rd**.
2. If necessary, open **Server Manager**.
3. In Server Manager, select **Add roles and features**, and then select **Next**.
4. Choose **Role-based** or **feature-based** installation, and then select **Next**.
5. On the **Server Selection** page, select **Next**.
6. On the **Server Roles** page, select the **Remote Access** check box, and then select **Next**.
7. On the **Features** page, select **Next**.

8. Select **Next**, and then on the **Role Service** page, select the **Routing** check box, select **Add Features**, and then select **Next**.
9. Select **Install**.
10. When installation completes, select **Close**.

Enable routing

1. In Server Manager, select **Tools**, and then select **Routing and Remote Access**.
2. In the **Routing and Remote Access** console, right-click or access the shortcut menu of the local server, select **Configure and Enable Routing and Remote Access**, and then select **Next**.
3. In the **Routing and Remote Access Server Setup Wizard**, on the **Configuration** page, select **Custom configuration**, and then select **Next**.
4. On the **Custom configuration** page, select the **LAN routing** check box, and then select **Next**.
5. When prompted, select **Finish**.
6. In the **Routing and Remote Access** dialog box, select **OK**.
7. When prompted, select **Start service**.

Configure routing

1. In the **Routing and Remote Access** console, expand **IPv4**.
2. Select **General**, and then right-click or access the shortcut menu of **General**.
3. Select **New Routing Protocol**.
4. In the **New Routing Protocol** dialog box, select **RIP Version 2** for **Internet Protocol**, and then select **OK**.
5. In the navigation pane, select **RIP**.
6. Right-click or access the shortcut menu of **RIP**, and then select **New Interface**.

7. Select the first network interface in the list, and then select **OK**.
8. In the **RIP Properties – *adapter name* Properties** dialog box, select the **Outgoing packet protocol** list. Notice that you can choose between the following protocols:
   - RIP version 1 broadcast
   - RIP version 2 broadcast
   - RIP version 2 multicast
   - Silent RIP

   Select **RIP version 2 multicast**.
9. Select the **Incoming packet protocol** list, select **RIP version 2 only**, and then select **OK**.
10. Right-click or access the shortcut menu of **RIP**, and then select **New Interface**.
11. Select the remaining network interface, and then select **OK**.
12. In the **RIP Properties – *adapter name* Properties** dialog box, select the **Outgoing packet protocol list**, and then select **RIP version 2 multicast**.
13. Select the **Incoming packet protocol list**, select **RIP version 2 only**, and then select **OK**.

Observe static routes

1. In the **Routing and Remote Access** console, select **Static Routes**.
2. Right-click or access the shortcut menu of **Static Routes**, and then select **Show IP Routing Table**.
3. Close the routing table.

Note

There are no further routes to observe because although RIP is installed, there are no other RIP routers to configure. Leave any VMs running for the next demonstration.

What is NAT?

Computers and devices that connect to the internet must be configured with public IP addresses. However, the number of public IPv4 addresses is becoming more limited, and organizations can't obtain a public IPv4 address for every organizational computer. Therefore, organizations use private IPv4 addressing for organizational computers.

Because private IPv4 addresses aren't routable on the internet, computers that are configured with private IPv4 addresses can't access the internet. By using NAT, organizations must obtain only one public IPv4 address to access the internet. NAT then translates the private IPv4 address into a public IPv4 address, which then provides internet access to organizational computers.

In Windows Server, a NAT server has two network adapters. One of these network adapters is configured with a private IPv4 address and connects to the organization's network, whereas the other network adapter is configured with a public IPv4 address and connects to the internet.

How does NAT work?

To connect a client computer to the internet by using NAT, you must configure the computer to use the NAT server as a default gateway. When a client computer on a private network requests access to a computer that's located on the internet—such as a web server—the NAT-enabled server translates the outgoing packets and then sends them to the web server on the internet. The NAT server also translates the response from the web server on the internet and then returns it to the client on the organization's network, as Figure 11 depicts.

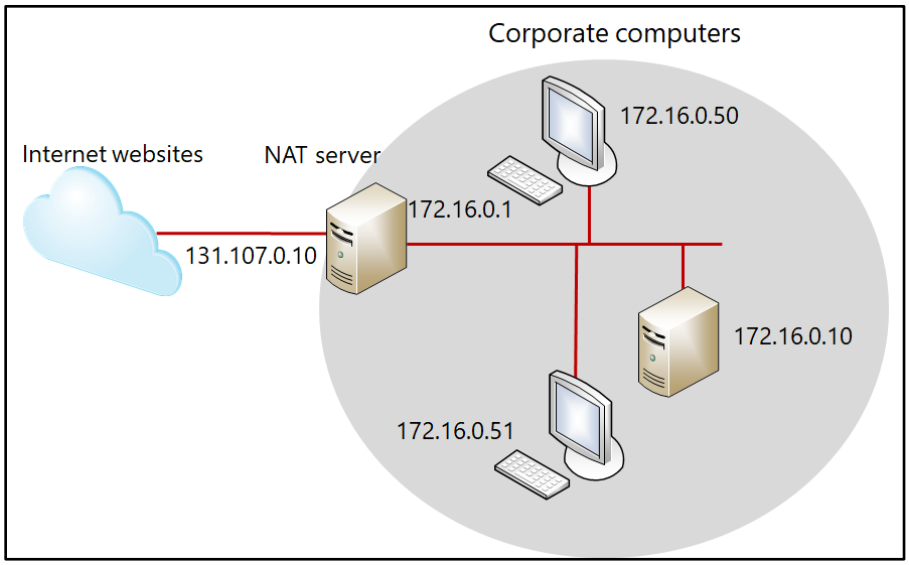

Figure 11. Implementing NAT

A NAT server secures an organizational network by hiding the IP addresses of the computers on that network. When a computer on the organization's network communicates with a web server on the internet, only the external IP address of the NAT server is visible to the internet web server. Furthermore, you can configure Windows Firewall with Advanced Security on the NAT server to protect your organization's network from internet security threats.

Demonstration: Implementing NAT in Windows Server

Your instructor will now demonstrate how to deploy and configure NAT in Windows Server.

Remove the router configuration

1. On the Windows Server computer on which you installed the Routing role, switch to the **Routing and Remote Access** console.

2. Right-click or access the shortcut menu of the local server, choose **Disable Routing and Remote Access**, and then select **Yes**.

Enable NAT

1. In the **Routing and Remote Access** console, right-click or access the shortcut menu of the local server, select **Configure and Enable Routing and Remote Access**, and then select **Next**.
2. In the **Routing and Remote Access Server Setup Wizard**, on the **Configuration** page, select **Network address translation (NAT)**, and then select **Next**.
3. On the **NAT Internet Connection** page, choose the network adapter with the public IPv4 address (**131.107.0.100**).
4. Select **Next**, and then select **Finish**.

Configure NAT

1. In the **Routing and Remote Access** console, expand **IPv4**.
2. Select **NAT**, right-click or access the shortcut menu of **NAT**, and then select **Properties**.
3. Select the **Address Assignment** tab.
4. Select **Automatically assign IP addresses by using the DHCP allocator** check box.
5. Select the **Name Resolution** tab.
6. Select the **Clients using Domain Name System (DNS)** check box, and then select **OK**.
7. In the **NAT node** details pane, right-click or access the shortcut menu of the network interface that's the public interface, and then select **Properties**.
8. In the **Properties** dialog box, select the **Services and Ports** tab. Notice the services that can be redirected.
9. Select **Cancel**.

Learning in action: Extending your network

Scenario

Lucerne Publishing has several departments in its head offices in Kensington, London. The haphazard way in which the network has grown with the organization has meant that departments are spread across all floors. The management team doesn't really want to restructure the cabling in the building to address these inconsistencies, nor does it want to move people around the building to match the cabling layout.

Questions

1. **What device could you install that would enable you to group computers together into logical networks even though those devices are distributed across all floors of the building?**

   A. Hubs

   B. Switches

   C. Routers

   D. VLANs

2. **You have set up two routers to manage traffic between three networks in the publishing department, as Figure 12 depicts. Network 1 has the subnet address 172.16.8.0/21, Network 2 has 172.16.16.0/21, and Network 3 has 172.16.24.0/21. A router connects Network 1 and 2. Another router connects Network 2 and 3.**

   Router 1 has two interfaces: 172.16.8.1/21 and 172.16.16.1/21. Router 2 has two interfaces: 172.16.16.2/21 and 172.16.24.1/21. Two routing tables (one for each router) contain incomplete entries that describe each interface and its connected networks.

   The routing table for Router 1 and Router 2 currently reads:

   172.16.8.0 172.16._._

   172.16.16.0 172.16._._

   172.16.24.0 172.16._._

   You must now complete the routing tables for Router 1. Select the correct configuration for routing table A's entry for network 172.16.24.0.

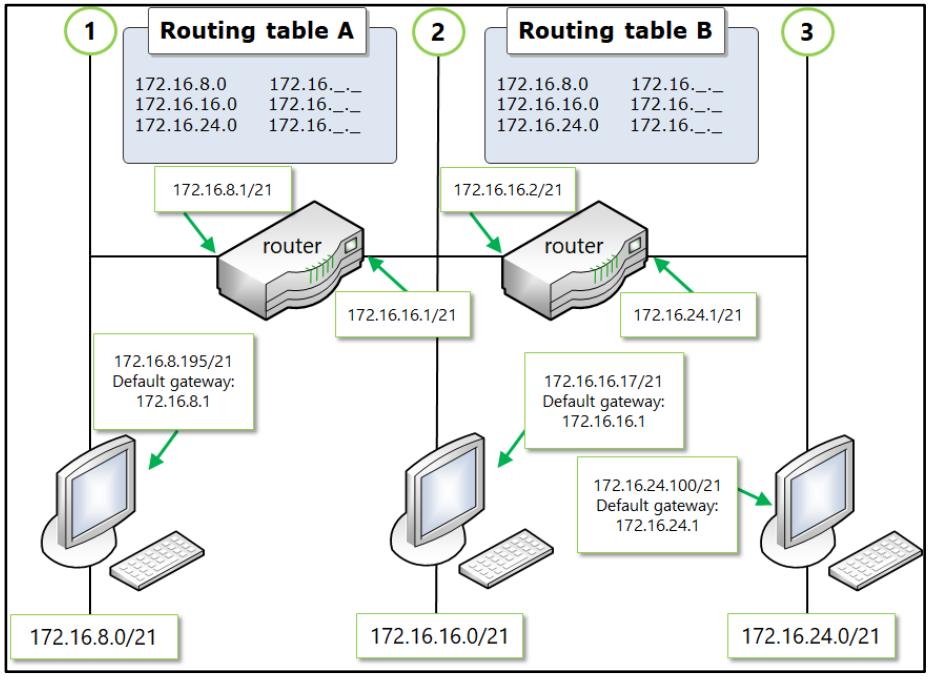

Figure 12. Configuration with two routers, three networks

Which of the following ranges would be correct?

A. 172.16.24.0 > 172.16.16.2

B. 172.16.24.0 > 172.16.24.1

C. 172.16.24.0 > 172.16.8.1

3. **Rather than configure manual routes for routers in the organization, you decide to deploy a routing protocol for your internal routers. Which of the following routing protocols would be most appropriate?**

A. RIP

B. EIGRP

C. BGP

D. OSPF

4. **You set up two routers to manage traffic between three networks in the publishing department, as Figure 13 depicts. For the network topology, Router 1 has two interfaces: 172.16.8.1/21 and 172.16.16.1/21. Router 2 has two interfaces: 172.16.16.2/21 and 172.16.24.1/21. After entering the routing tables, you discover that a device with an IP address of 172.16.16.17/20 in Network 2 can't communicate with other devices.**

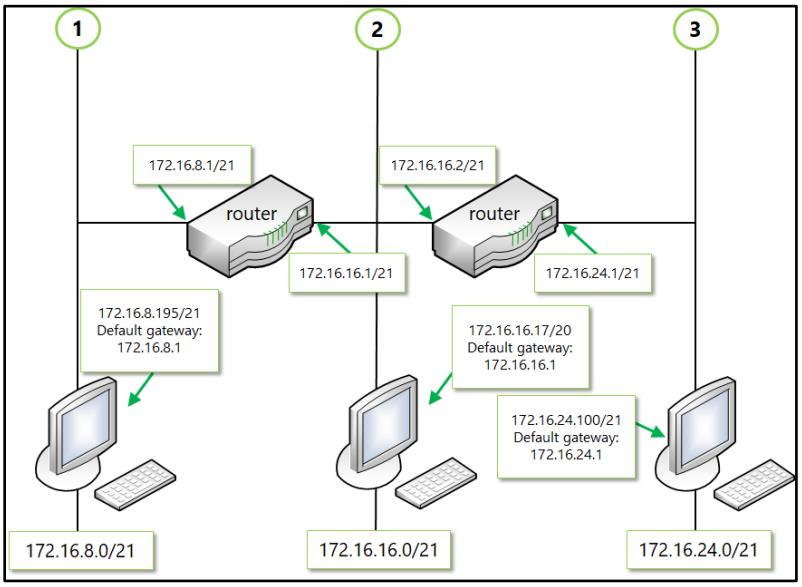

Figure 13. Continued depiction of configuration with two routers and three networks

After reviewing the network topology, what's the configuration error?

A. The router interface for the host is incorrect.

B. The IP address for the host is incorrect.

C. The subnet mask for the host is incorrect.

Test your knowledge

1. **Which network management protocol is often used to manage switches?**

   A. TCP/IP

   B. IGMP

   C. SNMP

   D. RIP

2. **Which of the following statements about routers is true? Choose all that apply.**

   A. A router operates at the data-link layer.

   B. A router is protocol dependent.

   C. A router is addressed by the hosts on the network.

3. **With routing, when a packet is determined to require remote delivery to a target host, the MAC address of which device is required by the local sending host?**

   A. Its own MAC address

   B. The MAC address of the default gateway of the target host

   C. The MAC address of its own default gateway

   D. The MAC address of the target host

Fill in the blanks for the following statements.

4. **Ethernet is an example of a ( ) network in which devices compete for bandwidth.**
5. **MAC tables in switches are also known as ( ) tables.**
6. **You can use the ( ) command to add a static route to a Windows computer.**
7. **True or false? The route 127.0.0.1 is a local loopback.**

   True

   False

8. **True or false? A layer 3 switch operates much like a network bridge.**

   True

   False

Study the scenario and answer the question.

9. **You're working for Lucerne Publishing in the IT planning department. Because of an increase in the number of connected devices and the commensurate increase in network traffic, you've been tasked with coming up with a plan to help manage the growth of the network. You measure network traffic in the research department LANs and realize that there are excessive network collisions. As a result, bandwidth utilization is being severely affected.**

   What device or devices should you consider deploying?

Study the scenario and answer the question.

10. **You're working for Lucerne Publishing in the IT planning department. You must come up with a plan to help manage the growth of the network, an increase in the number of connected devices, and the consequent increase in network traffic. You measure network traffic in the research department LANs and discover excessive network collisions. Closer inspection indicates that the network traffic on the research department LANs isn't related to research department computers.**

    How could implementing VLANs help?

Glossary

Term	Definition
Bridge	A layer 2 internetwork device that forwards frames to the appropriate network interface. It also forwards all broadcasts and multicasts.
Router	A layer 3 internetwork device that routes packets to appropriate subnets based on routes.
Switch	A device that can operate at layer 2, 3, or 4 to direct (switch) frames, packets, or datagrams, depending on level of operation, to appropriate connected networks.