CHAPTER 3

Planning, Implementing, and Maintaining a Network Infrastructure

Terms you'll need to understand:

- ✓ Internet Group Management Protocol (IGMP)
- ✓ Internet Control Message Protocol (ICMP)
- ✓ Quality of Service (QoS)
- ✓ Open Systems Interconnect (OSI)
- ✓ Internet Protocol (IP)
- ✓ Address Resolution Protocol (ARP)
- ✓ Transmission Control Protocol (TCP)
- ✓ User Datagram Protocol (UDP)
- ✓ Subnet mask
- ✓ Classless Inter-Domain Routing (CIDR)
- ✓ Packet Internet Groper (PING)
- ✓ `IPCONFIG` command
- ✓ `NBTSTAT` command
- ✓ `NETSTAT` command
- ✓ `ROUTE` command
- ✓ `HOSTNAME` command
- ✓ `TRACERT` command
- ✓ `PATHPING` command
- ✓ File Transfer Protocol (FTP)
- ✓ `TELNET` command
- ✓ Remote Copy Protocol (RCP)
- ✓ Remote Shell (RSH)
- ✓ Remote Exec (REXEC)
- ✓ Dynamic Host Configuration Protocol (DHCP)
- ✓ Automatic Private IP Addressing (APIPA)
- ✓ DHCP Relay Agent
- ✓ Domain Name System (DNS)
- ✓ Time to Live (TTL)
- ✓ Recursive
- ✓ Iterative
- ✓ Incremental zone transfer (IXFR)
- ✓ Full zone transfer (AXFR)
- ✓ (A) records
- ✓ CNAME (canonical name) records
- ✓ MX (Mail Exchanger) records
- ✓ NS (Name Server) records
- ✓ PTR (Pointer) records
- ✓ SOA (Start of Authority) records
- ✓ SRV (Service) records
- ✓ Windows Internet Naming Service (WINS)

Techniques you'll need to master:

- ✓ Installing DNS, WINS, and DHCP
- ✓ Configuring clients to use DHCP, DNS, WINS, and APIPA
- ✓ Configuring clients to use Dynamic Update
- ✓ Configuring DHCP to perform Dynamic Update on behalf of clients
- ✓ Configuring DHCP scopes and optional parameters
- ✓ Configuring an IP subnet
- ✓ Analyzing IP addressing requirements

Transmission Control Protocol/Internet Protocol (TCP/IP) is a connection-oriented, Internet-standard, routable protocol in use on a majority of networks, including the Internet. The protocol suite supports connectivity across a number of dissimilar platforms and supports the main workload of most enterprises today that are designed in a client/server configuration.

Some subtle changes have been incorporated into the TCP/IP suite for Windows Server 2003. Internet Group Management Protocol (IGMP) version 3 adds support for source-based filtering and reporting while maintaining backward-compatibility with version 2. You can also use other settings so that systems can be configured to use an alternate, manually configured IP address instead of one that a Dynamic Host Configuration Protocol (DHCP) server provides. Autoconfiguration of the enabled network card interface (NIC) metric is also available; this feature determines the best routing metric for each interface's default gateway, based on its speed. Support for TCP/IP version 6 has also been added in Windows Server 2003.

Overview of TCP/IP

TCP/IP is a network communication protocol suite. It can be used as a communications protocol on private networks and is the default protocol in use on the Internet. When you set up any system to have direct access to the Internet, whether it is via dial-up or a high-speed technology, your system needs to use TCP/IP whether it is a Windows-based system or not.

Also, if systems need to communicate to other TCP/IP systems on the local area network (LAN) or wide area network (WAN), they often use TCP/IP as well.

NOTE

Indirectly connected computers, such as those on a LAN that connect to the Internet via certain default gateways, certain types of routers, proxy servers, or other indirect means, do not necessarily need to use TCP/IP. They need to use only the network protocol in use on the LAN, and that LAN protocol communicates with the directly connecting mechanism (default gateway, router, proxy server, or other direct device). That directly connected device needs to use the Internet default protocol of TCP/IP.

For Internet Security and Acceleration (ISA) servers, systems must use TCP/IP because it is the supported protocol for ISA.

OSI Model

TCP/IP is technically made up of two protocols. The upper layer, Transmission Control Protocol, is responsible for breaking data down into smaller packets to be transmitted over the network from a sending system (local and Internet), and

the TCP layer on the receiving system reassembles the packets it receives into the original data structure. The lower layer, Internet Protocol, addresses each packet so that it gets delivered to the correct remote system. Each routing device on the network, be it a hardware router or a server system performing routing functions, checks the destination address to see where to forward the message.

The TCP/IP protocol suite maps to a four-layer conceptual model, which parallels the seven-layer Open Systems Interconnect (OSI) protocol model described in the following list:

- *Physical layer*—This layer defines the interface between the network medium (such as ethernet or token ring) and the hardware device (such as a NIC). Multiplexers, hubs, and repeaters are just a few examples of the components found at this layer of the OSI model.
- *Data Link layer*—This layer is divided into two sublayers: Logical Link Control (LLC), which handles error correction and flow control, and Media Access Control (MAC), which handles communication with the NIC. Bridges and switches are components that operate at this layer of the OSI model.
- *Network layer*—This layer translates logical network address and names to MAC addresses for routing data packets over a network. A number of protocols run at the Network layer, including IP, Address Resolution Protocol (ARP), Reverse ARP (RARP), Internet Control Message Protocol (ICMP), Routing Information Protocol (RIP), Open Shortest Path First (OSPF), IGMP, Internetwork Packet Exchange (IPX), NWLink (the Microsoft version of the IPX/SPX protocol suite), and NetBIOS Enhanced User Interface (NetBEUI). Brouters, routers, and some types of ATM switches can be found at this layer of the OSI model.
- *Transport layer*—This layer provides an additional connection below the Session layer and assists with managing some data flow control between hosts. Data is divided into packets on the sending node, and the receiving node's Transport layer reassembles the message from packets. This layer is also responsible for error checking to guarantee error-free data delivery, and requests a retransmission if necessary. It is also responsible for sending acknowledgments of successful transmissions back to the sending host. A number of protocols run at the Transport layer, including TCP, ARP, RARP, Sequenced Packet Exchange (SPX), and NWLink. Gateways and certain types of routers can be found at this layer of the OSI model.

- *Session layer*—This layer establishes, maintains, and ends sessions between transmitting hosts and controls which host can transmit data at a given interval and for how long. A number of protocols run at the Session layer, including Named Pipes, NetBIOS Names, Remote Procedure Calls (RPC), and Mail Slots. Gateways and certain types of proxy servers operate at this layer of the OSI model.
- *Presentation layer*—This layer translates data from the way applications understand it to the way networks understand it. It is responsible for protocol conversions, data encryption and decryption, and data compression and decompression when the network is considered. Gateways and certain types of redirectors operate at this layer of the OSI model. There are no protocols that normally operate in this layer of the OSI model.
- *Application layer*—This layer allows access to network services for applications specifically written to run over the network. Some protocols found at this OSI layer include File Transfer Protocol (FTP), Trivial FTP (TFTP), Bootstrap Protocol (BOOTP), Simple Network Management Protocol (SNMP), Simple Mail Transfer Protocol (SMTP), Telnet, NetWare Core Protocol (NCP), and Server Message Block (SMB) .

The four-layer conceptual model for the TCP/IP protocol suite is as follows:

- *Network Interface layer*—This layer is responsible for putting bits on the wire and correlates closely with the OSI model's Physical layer and Data Link layer.
- *Internet layer*—This layer is responsible for encapsulating data packets into Internet datagrams. The Internet layer correlates, for the most part, with the OSI model's Network layer. Four Internet protocols operate at this layer:
  - IP supports connectionless packet delivery for all other protocols, such as TCP or User Datagram Protocol (UDP). IP does not guarantee packet arrival or correct packet sequence, nor does it acknowledge packet delivery. These tasks are left to the application using the network or higher-level protocols, such as TCP. IP is responsible for addressing and routing packets only; error correction is left to the application or to higher-level protocols.
  - ARP is responsible for mapping IP addresses to physical machine addresses called MAC addresses. IP broadcasts a special ARP inquiry packet containing the destination system's IP address, and that system replies by sending its physical address to the requester.

  - ICMP is charged with message control and error-reporting between network hosts. Higher-level protocols use this information to recover from transmission errors.
  - IGMP allows hosts to report their multicast group membership to multicast routers. With multicasting, hosts can send multicast traffic to a single MAC address, so multiple nodes can process the traffic.
- *Transport layer (also called Host-to-Host Transport)*—This layer basically (but not entirely) correlates with the OSI model's Transport layer. The two Transport layer protocols, TCP and UDP, provide communication sessions between systems.
  - TCP is a connection-oriented protocol that guarantees data delivery by assigning a sequence number to each transmitted data segment so that the receiving host can send an acknowledgment (ACK) to verify that the data was received intact. If an ACK is not received or there was a transmission error, the data is sent again.
  - UDP is a connectionless protocol that does not guarantee delivery or correct sequencing of packets. Applications that use UDP are typically tasked with the responsibility of ensuring data delivery because the protocol does not. UDP is often used instead of TCP because of its lower overhead. TFTP is an example of an application that uses UDP.
- *Application layer*—This layer is where network-aware applications operate. Network applications most commonly use two TCP/IP services, Winsock and the NetBT interface.

IP Addressing

IP version 4 (IPv4) addresses are made up of four 8-bit fields (*octets*)—32 bits total. There are five IPv4 address classes: A, B, C, D, and E.

IPv4 addresses consist of a network ID and a host ID. The *network ID* identifies the numeric network name of the physical network where the hosts exist. The *host ID* identifies the numeric network name of the individual TCP/IP host on a network. For example, in the Class A IP address 10.0.0.1, 10 represents the network ID and 0.0.1 represents the host ID. The numeric host ID must be unique on the internal network—that is, no two nodes on a network can have the same network ID and host ID. Using the previous example, only one host can be assigned the host ID of 0.0.1 on the given network.

NOTE

You can have two hosts with the same numeric IP hostname of 16.72.28 if one is on network 111 and another is on network 112. (The full IP addresses of these hosts would be 111.16.72.28 and 112.16.72.28. The subnet mask would be 255.0.0.0.)

A *subnet mask* is used to divide an entire TCP/IP address in an effort to define which part of the address is the network number and which part is the host system's numeric identifier. The bits in a subnet mask are set consecutively from left to right. For example, the subnet mask 255.128.0.0 is valid because all eight bits are set in the first two octets and the first bit of the next octet is also set (11111111.10000000.00000000.00000000). The subnet mask 255.64.0.0 is not valid because it has a "missing" bit, which is not allowed (11111111.01000000.00000000.00000000).

NOTE

Bit values are held to a specific order, from the Most Significant Bit (MSB) to the Least Significant Bit (LSB). From left to right, these designations are 128, 64, 32, 16, 8, 4, 2, and 1. Each bit that's set is noted by a "1" (showing that the bit is "on" or "enabled"), and bits are added together to give you the address. The IPv4 address 171.144.62.12 converts to a binary number of 10101011.10010000.00111110.00001100.

EXAM ALERT

You need to have a fairly good understanding of host IDs, network IDs, subnetting, and masks for just about any Microsoft certification exam. Any exams that introduce information about networking require you to have at least basic knowledge of TCP/IP addressing.

Subnet Masks

When assigning IP addresses, each host requires a subnet mask to determine which part of an IP address to use as the network ID and which to use as the host ID.

The default subnet masks for the three IP address classes are

- Class A - 255.0.0.0
- Class B - 255.255.0.0
- Class C - 255.255.255.0

For example, the default subnet mask for a Class C address is 255.255.255.0, which means the first three octets identify the network and the last octet indentifies the host.

The subnet mask is also used to determine whether the destination host is on the local subnet or a remote subnet. The subnet mask of the local host is compared against the IP address of the destination host and, through a process known as *anding*, it is determined whether the destination IP address is the local or a remote network. If the destination IP address within a packet is on a remote network, the packet is sent to the default gateway.

Basically, the number of 1's in the binary address of the subnet mask are masked against the IP address to determine if the address is on the local network or a remote network. When the bits of the subnet mask are compared against the bits in the IP address, all combinations of 1's and 0's result in a value of 0, except for 1 and 1, which results in a value of 1.

Let's take at an example of how this process works. The source host has an IP address of 192.168.0.10 and a subnet mask of 255.255.255.0. The destination host has an IP address of 192.168.20.2.

IP address 11000000 10101000 00000000 00001010 (192.168.0.10)

Subnet mask 11111111 11111111 11111111 00000000 (255.255.255.0)

Results 11000000 10101000 00000000 00000000

IP address 11000000 10101000 00010010 00000010 (192.168.20.20)

Subnet mask 11111111 11111111 11111111 00000000 (255.255.255.0)

Results 11000000 10101000 00010010 00000000

As you can see from the preceding example, the source IP address is anded against the subnet mask. The destination address is anded against the subnet mask assigned to the source host. If the results are not the same, the destination host is on a different network or subnet. Conversely, if the results are the same, it is determined that the destination host is on the local network.

The original IP definitions set five classes of IP addresses, from A through E. (A, B, and C are for general-purpose use, D is used for multicasting, and E is reserved.) These classes made it possible to use one portion of the 32-bit IP address scheme for the network address and the remaining portion for nodes on the network.

In the past, some networks needed more addresses for systems than the 254 a Class C address supplies. This was a major contribution to the shortage of IP addresses. Organizations often requested a Class B range that offered 65,534

available addresses rather than a few Class C ranges that might have suited their needs. The result was that many addresses within their allotted Class B blocks went unused.

However, Classless Inter-Domain Routing (CIDR) addressing is now used more often for IPv4 addressing schemes. It effectively "removes" the class from an address for the purpose of combining ranges, so it makes better use of the limited number of remaining available IPv4 addresses. A CIDR network address looks like this:

```
222.175.14.00/18
```

The network address is 222.175.14.00. The `/18` specifies that the first 18 bits of the address are the network part of the address, which leaves the last 14 bits for the network hosts' address.

Both Border Gateway Protocol (BGP) and OSPF support CIDR. Older gateway protocols, such as Exterior Gateway Protocol (EGP) and Routing Information Protocol version 1 (RIPv1), do not support CIDR. Because CIDR supports multiple subnet masks per subnet, it requires routers that support more advanced interior routing protocols, such as RIPv2 and OSPF.

NOTE

A, B, and C class networks support a single subnet mask and can use RIPv1.

Create an IP Subnet Scheme

Implementing subnets helps control network traffic and enables network administrators to create smaller collision domains. Every node on the same physical ethernet network sees all data packets sent out on the network, which results in multiple collisions and affects network performance. Routers or gateways separate networks into subnets. Subnet masks on each node allow nodes on the same subnetwork to continue communicating with one another and with the routers or gateways they use to send their messages.

Subnet masking enables you to identify the network ID and host (node) ID of an IP address. The following example is a default Class B subnet mask:

```
10110110.10100101.00110111.01100010 182.165.55.98
11111111.11111111.00000000.00000000 255.255.000.000
-----------------------------------------------------
10110110.10100101.00000000.00000000 182.165.000.000
```

```
IP Address       : 182.165.55.98
Address Class    : B
Network Address  : 182.165.0.0

Subnet Address   : 182.165.48.0
Subnet Mask      : 255.255.240.0
Subnet bit mask  : 11111111.11111111.11110000.00000000
Subnet Bits      : 20
Host Bits        : 12
Possible Number of Subnets : 16
Hosts per Subnet : 4094

Selected Subnet  : 182.165.0.0/255.255.240.0
Usable Addresses : 4094
     Host range  : 182.165.0.1  to  182.165.15.254
     Broadcast   : 182.165.15.255
```

To subnet networks further, more bits can be added to the subnet mask for a class of addresses.

The following example is a Class B address using an additional bit subnet mask of 240. Notice that instead of having the single subnet and 65,534 hosts per subnet allowed under the default subnet mask, you can have up to 16 subnets with up to 4,094 hosts per subnet by using a subnet mask of 255.255.240.000 (Table 3.1 shows a sample IP addressing scheme):

```
10110110.10100101.00110111.01100010 182.165.55.98
11111111.11111111.11110000.00000000 255.255.240.000 Subnet Mask
------------------------------------------------------
IP Address       : 182.165.55.98
Address Class    : B
Network Address  : 182.165.0.0

Subnet Address   : 182.165.48.0
Subnet Mask      : 255.255.240.0
Subnet bit mask  : 11111111.11111111.11110000.00000000
Subnet Bits      : 20
Host Bits        : 12
Possible Number of Subnets : 16
Hosts per Subnet : 4094

Selected Subnet  : 182.165.0.0/255.255.240.0
Usable Addresses : 4094
     Host range  : 182.165.0.1  to  182.165.15.254
     Broadcast   : 182.165.15.255
```

TABLE 3.1 Example of an IP Addressing Scheme

Subnet	Mask	Subnet Size	Host Range	Broadcast
182.165.0.0	255.255.240.0	4094	182.165.0.1 to 182.165.15.254	182.165.15.255
182.165.16.0	255.255.240.0	4094	182.165.16.1 to 182.165.31.254	182.165.31.255
182.165.32.0	255.255.240.0	4094	182.165.32.1 to 182.165.47.254	182.165.47.255
182.165.48.0	255.255.240.0	4094	182.165.48.1 to 182.165.63.254	182.165.63.255
182.165.64.0	255.255.240.0	4094	182.165.64.1 to 182.165.79.254	182.165.79.255
182.165.80.0	255.255.240.0	4094	182.165.80.1 to 182.165.95.254	182.165.95.255
182.165.96.0	255.255.240.0	4094	182.165.96.1 to 182.165.111.254	182.165.111.255
182.165.112.0	255.255.240.0	4094	182.165.112.1 to 182.165.127.254	182.165.127.255
182.165.128.0	255.255.240.0	4094	182.165.128.1 to 182.165.143.254	182.165.143.255
182.165.144.0	255.255.240.0	4094	182.165.144.1 to 182.165.159.254	182.165.159.255
182.165.160.0	255.255.240.0	4094	182.165.160.1 to 182.165.175.254	182.165.175.255
182.165.176.0	255.255.240.0	4094	182.165.176.1 to 182.165.191.254	182.165.191.255
182.165.192.0	255.255.240.0	4094	182.165.192.1 to 182.165.207.254	182.165.207.255
182.165.208.0	255.255.240.0	4094	182.165.208.1 to 182.165.223.254	182.165.223.255
182.165.224.0	255.255.240.0	4094	182.165.224.1 to 182.165.239.254	182.165.239.255
182.165.240.0	255.255.240.0	4094	182.165.240.1 to 182.165.255.254	182.165.255.255

When you use standard subnet masks in classful IP addressing schemes, you can plan how many hosts you can support per subnet and how many subnets are available for use. Table 3.2 shows classful IP addressing schemes and uses 255.x.0.0 as the default mask for Class A addresses, 255.255.x.0 as the default mask for Class B class addresses, and 255.255.255.x as the mask for Class C addresses. In these classes, the *X* is the subnet mask variable in the table's Subnet Mask column. The table identifies how many subnets ID are supported by each subnet mask and the maximum number of hosts per subnet.

TABLE 3.2 Subnet Masking for Classful IP Addressing

Subnet Mask	Number of Subnets in Classful Range	Number of Class A Hosts per Subnet	Number of Class B Hosts per Subnet	Number of Class C Hosts per Subnet
0	1	16,777,214	65,534	254
128	2	8,388,606	32,766	126
192	4	4,194,302	16,382	62
224	8	2,097,150	8,190	30
240	16	1,048,574	4,094	14

TABLE 3.2 *Continued*

Subnet Mask	Number of Subnets in Classful Range	Number of Class A Hosts per Subnet	Number of Class B Hosts per Subnet	Number of Class C Hosts per Subnet
248	32	524,286	2,046	6
252	64	262,142	1,022	2
254	128	131,070	510	N/A
255	256	65,534	254	N/A

IP Address Classes

IP addresses are organized into different address classes that define the number of bits out of the 32 that are used to identify the network and which are used to identify hosts on a network. By examining the address classes, you can also determine the number of networks and the number of hosts.

TCP/IP Class A Addresses

Class A addresses have an official start address of 0.0.0.0 and an official ending address of 127.255.255.255. However, the last usable client address in the range is 126.255.255.254 because the 127.*x.x.x* range is used for internal host loopback.

The full range of addresses that can be assigned to hosts is 1.0.0.1 to 126.255.255.254, with 126.255.255.255 as the broadcast address. The local host uses 0.0.0.0 when it has been configured to use a DHCP server but cannot reach one and cannot assign itself an address using APIPA. (This situation would be unusual.)

There are 126 Class A networks total, and each is allowed to have up to 16,777,214 hosts.

Three IP network addresses are reserved for private networks as defined in Request for Comment (RFC) 1918. The Class A range is 10.0.0.0 to 10.255.255.255, with a subnet mask of 255.0.0.0.

These addresses can be used by anyone setting up internal IP networks, such as a lab or home LAN behind a Network Address Translation (NAT) server, proxy server, or router. It is always safe to use them because routers on the Internet never forward packets coming from these addresses.

TCP/IP Class B Addresses

The Class B range of IP addresses starts with address 128.0.0.0 and ends at address 191.255.255.255. IP addresses 128.0.0.1 to 191.255.255.254 are the usable range of Class B addresses for node assignment.

Three IP network addresses are reserved for private networks, as defined in RFC 1918. The Class B range is 172.16.0.0 to 172.31.255.255, with the subnet mask 255.240.0.0. These addresses can be used by anyone setting up internal IP networks, such as a lab or home LAN behind a NAT server, proxy server, or router. It is always safe to use these addresses because routers on the Internet never forward packets coming from these addresses.

TCP/IP Class C Addresses

The Class C range of IP addresses starts at address 192.0.0.0 and ends at 223.255.255.255. IP addresses 192.0.0.1 to 223.255.255.254 are the usable range of Class C addresses for node assignment.

Three IP network addresses are reserved for private networks, as defined in RFC 1918. The Class C range is 192.168.0.0 to 192.168.255.255, with the subnet mask 255.255.0.0. These addresses can be used by anyone setting up internal IP networks, such as a lab or home LAN behind a NAT server, proxy server, or router. It is always safe to use them because routers on the Internet never forward packets coming from these addresses.

TCP/IP Class D Addresses

The Class D IP addresses range from 224.0.0.0 through 239.255.255.255. Internet Assigned Numbers Authority (IANA) has set aside this range as a special class of addresses for multicast uses. ISPs are unable to allocate Class D address space to their customers because IANA is the only body through which these addresses can be allocated.

Allocation of Class D addresses is required only if you want to be a multicast source. You can still receive multicast data without needing a separate Class D address.

TCP/IP Class E Addresses

IANA has set aside Class E IP addresses from 240.0.0.0 to 254.255.255.255 as a special class of addresses for experimental and future use. The IP address 255.255.255.255 broadcasts to all hosts on the local network and, therefore, is not considered part of the Class E IP addresses.

Well-Known Ports

A number of well-known ports (0–1023) are used by different services on computers. For a single IP address on one system to offer all possible services to a network, each service must function on its own TCP or UDP port from that IP address.

You can find a helpful table at http://www.networksorcery.com that includes links to definitions and additional notes for some services. The following ports and associated protocols are the most important ones to remember for the certification exam:

- *20*—FTP—data
- *21*—FTP—control
- *22*—Secure Shell (SSH)
- *23*—Telnet
- *25*—SMTP
- *37*—Time Protocol (Time)
- *49*—Terminal Access Controller Access Control System (TACACS), TACACS+
- *53*—DNS
- *67*—BOOTP—server
- *68*—BOOTP—client
- *69*—TFTP
- *70*—Gopher
- *79*—Finger
- *80*—Hypertext Transfer Protocol (HTTP)
- *88*—Kerberos
- *109*—Post Office Protocol version 2 (POP2)
- *110*—Post Office Protocol version 3 (POP3)
- *115*—Simple File Transfer Protocol (SFTP)
- *119*—Network News Transfer Protocol (NNTP)
- *123*—Network Time Protocol (NTP)
- *137*—NetBIOS Name Service
- *138*—NetBIOS Datagram Service
- *139*—NetBIOS Session Service
- *143*—Internet Message Access Protocol (IMAP)
- *153*—Simple Gateway Monitoring Protocol (SGMP)
- *161*—SNMP

- *162*—SNMP—traps
- *179*—BGP
- *389*—Lightweight Directory Access Protocol (LDAP), Connection-less Lightweight X.500 Directory Access Protocol (CLDAP)
- *443*—HTTP over Secure Socket Layer/Transport Layer Security (SSL/TLS)—HTTPS
- *464*—Kerberos change/set password
- *500*—ISAKMP, Internet Key Exchange (IKE)
- *546*—DHCPv6 client
- *547*—DHCPv6 server
- *631*—Internet Printing Protocol (IPP)

Plan a TCP/IP Network Infrastructure Strategy

One of the major objectives on exam 70-293 is Planning a TCP/IP Network Infrastructure Strategy. The following section is dedicated to discussing the important topics that fall under this objective.

Analyze IP Addressing Requirements

Before you can successfully and effectively implement a TCP/IP network infrastructure, you have to identify your IP addressing requirements. Some of the things that you need to consider include

- Will you require a public or private IP address range on the internal network?
- How will client computers be assigned IP addresses?
- Does the network require multiple subnets?

Public/Private IP Address

Three ranges of IP addresses are reserved, meaning they are not valid on the Internet. Therefore, you can use one of these private ranges on your private network. Of course, one of the disadvantages to this is that a proxy server or NAT

server is needed for Internet access because the private IP address must be mapped to a public one. In terms of advantages, private IP addressing is more cost effective, can accommodate growth on your network, and can increase security.

NOTE

You can read more about the IP address ranges reserved for private networks in RFC 1918 (http://www.faqs.org/rfcs/rfc1918.html).

If you do decide to implement a private IP address range, you can use IP addresses from any of the following classes:

Class A 10.0.0.0–10.255.255.255

Class B 172.16.0.0–172.31.255.255

Class C 192.16.8.0.0–192.168.255.255

One of the decisions you are faced with when designing a TCP/IP network is whether you want to use a private IP address range or public IP addresses. Of course, there are disadvantages and advantages to both of them. In making your decisions, keep in mind that any computers that have a direct connection to the Internet will require at least one public IP address. However, for those computers with no direct Internet connection, you have the option of using public or private addresses.

IP Address Assignment

IP addresses can be dynamic or static. With static IP addressing, a computer (or other device) always uses the same IP address. With dynamic addressing, the IP address changes.

There are several ways of configuring a host with an IP address. You can do so manually or automatically, using a service such as DHCP or Automatic Private IP Addressing (APIPA).

Manual IP Addressing

For manual IP address assignment, an administrator or similarly delegated person manually enters a static IP address and other information, such as the subnet mask and default gateway, DNS server, WINS server, and so forth.

Dynamic Host Configuration Protocol

DHCP dynamically assigns randomly available IP addresses from available scopes to DHCP clients. This allows administrators to automatically assign IP addresses to clients without actually having to set all the parameters (default gateway, DNS servers, and so on) for each system, as with manual IP address assignment.

The DHCP lease process begins over UDP ports 67 and 68 as a broadcast message from the client system. For DHCP clients to contact DHCP servers on remote networks successfully, the IP routers must be RFC 1542–compliant. These routers support forwarding DHCP broadcasts off the local subnet. If the routers are not compliant, a DHCP Relay Agent must be in use on that subnet.

The DHCP Relay Agent is available through the Routing and Remote Access MMC on Windows Server 2003 (see Figure 3.1). Systems configured in the role of a DHCP server should not be configured as DHCP Relay Agents because both services use UDP ports 67 and 68 and degrade each other's services if they are installed together. On a single subnet, there is no practical need to do this because the DHCP server should simply respond to user requests on the subnet.

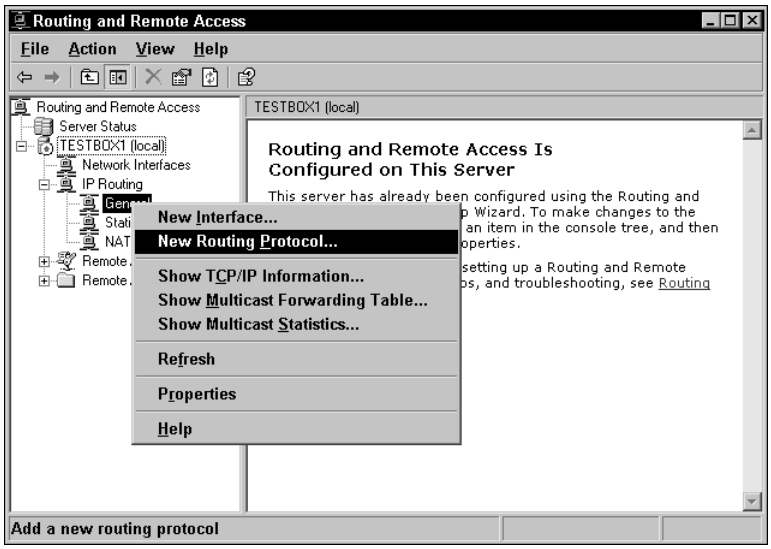

FIGURE 3.1 The first step in configuring the DHCP Relay Agent service in the Routing and Remote Access MMC.

Figure 3.2 shows the available routing protocols in the New Routing Protocol dialog box: DHCP Relay Agent, IGMP Router and Proxy, Open Shortest Path First (OSPF), and RIP Version 2 for Internet Protocol.

When a client first starts, it sends out a DHCPDISCOVER broadcast message to all addresses (255.255.255.255). The message contains the client's hardware (MAC) address and hostname. (The client also sends this message when its original lease has expired and cannot be renewed.) All available DHCP servers configured to respond receive the DHCPDISCOVER broadcast and send a DHCPOFFER broadcast message back with the following information:

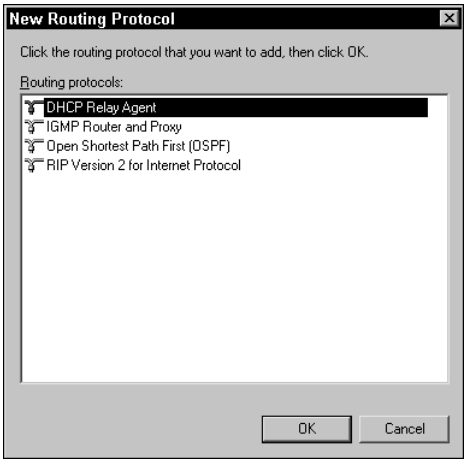

FIGURE 3.2 You can add routing protocols through the New Routing Protocol dialog box.

- The client's hardware address
- An offered IP address
- Subnet mask
- Length of the lease
- A server identifier (the IP address of the offering DHCP server)

The DHCP client selects the IP address from the first offer it receives and responds with a DHCPREQUEST broadcast message that includes the IP address of the server whose offer was accepted. All the other DHCP servers then retract their lease offers and mark those addresses available for the next IP lease request.

The DHCP server whose lease was accepted responds with a DHCPACK broadcast message, which contains the valid lease period for that IP address and other configuration information outlined in the scope, such as router information (default gateway), subnet mask, and so forth. After the DHCP client receives this acknowledgment, TCP/IP is completely initialized, and the client can use the IP address for communication.

Automatic Private IP Addressing

When a DHCP client sends out the DHCPDISCOVER broadcast message, it waits 1 second for an offer. If the client does not receive a response from a DHCP server, it rebroadcasts the request three times at 9-, 13-, and 16-second intervals, with a random offset length of 0ms and 1000ms. If an offer is not received after the four requests, the client retries once every five minutes.

Beginning with Windows 98, DHCP clients can configure themselves by using APIPA and the DHCP client service. After the four attempts to receive an IP address have failed, the DHCP client auto-configures its IP address and subnet mask using the reserved Class B network address 169.254.0.0 and the subnet mask 255.255.0.0. No default gateway is used, so systems that use APIPA are not routable.

The DHCP client tests for an address conflict to make sure that the IP address it has chosen is not already in use on the network. To do this, it broadcasts its selection from the range to the local subnet. If a conflict is found, the client selects another IP address and continues this process up to 10 attempts.

After the DHCP client makes sure that the address it has chosen is not in conflict with another system on the subnet, it configures its network interface with the IP address. The client then continues to check for a response from the DHCP server every five minutes. If a DHCP server becomes available, the client drops its APIPA settings and uses the address the DHCP server offers at that time.

Optimize TCP/IP Performance

So far you have learned about subnetting and configuring network systems in address class ranges in an effort to optimize TCP/IP configuration, but some other points should be mentioned as well. You need to be sure, above all else, that you understand your network configuration and behavior. Although you can take a few steps to fine-tune TCP/IP traffic, network topology plays a big role.

For TCP/IP specifically, there is the TCP/IP Receive Window Size setting, which is the buffer threshold for inbound packets. The default setting for ethernet networks is 17,520 bytes; when this threshold is met, the receiving system sends out an acknowledgement that the data has been received. This process of sending and acknowledging during a data transmission session repeats every 17,520 bytes until all data has been transmitted. As an administrator, you can adjust this acknowledgement setting to optimize transmissions.

Other settings on the network's Physical and Data Link layers are beyond normal administrative control. Maximum Transmission Units (MTUs), for example, are based on the type of network that is installed. For example, 16Mbps token-ring networks have an MTU setting of 17,914 bytes; 4Mbps token-ring networks have an MTU setting of 4,464 bytes. Ethernet deployments are limited to a 1,500 byte MTU setting. As an analogy, think of the MTU as an envelope in which data is carried.

The Maximum Segment Size (MSS) setting determines the largest segment that can be carried in the MTU. (Think of it as the pages of a letter in an envelope.) This setting also varies depending on the framework. Obviously, the MSS for token-ring networks will be larger than the MSS for ethernet networks.

Networks must consider application requirements when implementing certain services and protocols to optimize bandwidth. Quality of service (QoS) can also be implemented on networks to optimize bandwidth. The main issue on most networks is that all the associated networking equipment needs to support the Resource Reservation Protocol (RSVP). Networks also have certain application requirements to consider, such as the following:

- Routers forward traffic on a best-effort basis as they receive it. Video conferencing and streaming media suffer when available bandwidth is low.
- QoS Admission Control Service (QoS ACS) handles bandwidth on a subnet-to-subnet basis.
- Subnet Bandwidth Management (SBM) manages the use of network resources on a subnet.
- RSVP is a signaling protocol that enables sender and receiver systems to set up a reserved QoS session. RSVP messages carry the reservation request in an effort to maintain the QoS session. This is why each router and switch along the communication path between the sender and receiver needs to support RSVP.
- Traffic Control uses the packet classifier to separate packets into queues based on their priority. The Packet Scheduler manages the queues set up by the packet classifier.

Troubleshooting TCP/IP Addressing Problems

Windows XP Professional and Windows Server 2003 offer several native programs that an administrator can use to troubleshoot TCP/IP issues. Some are full-fledged tools in their own right, such as FTP, but they can help in determining what might be affecting a TCP/IP network. Many of these TCP/IP troubleshooting tools are discussed in the sections that follow.

The PING Command

The PING command can be used to test network connectivity from a local system by sending an ICMP message to a remote host or gateway. On external networks such as the Internet, the use of PING might be somewhat limited, depending on how routers and firewalls are configured; many do not allow ICMP traffic. If the remote host receives the message, it responds with a reply message. PING notes the IP address, the number of bytes in the message, how long it took to reply (in milliseconds [ms]), and the length of Time –to Live (TTL) in seconds and shows any packet loss in terms of percentages, as shown here:

```
D:\>ping 192.168.1.225
Pinging 192.168.1.225 with 32 bytes of data:
Reply from 192.168.1.225: bytes=32 time<10ms TTL=128
Reply from 192.168.1.225: bytes=32 time<10ms TTL=128
Reply from 192.168.1.225: bytes=32 time<10ms TTL=128
Reply from 192.168.1.225: bytes=32 time<10ms TTL=128
Ping statistics for 192.168.1.225:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milliseconds:
    Minimum = 0ms, Maximum =  0ms, Average =  0ms
Usage: ping [-t] [-a] [-n count] [-l size] [-f] [-i TTL] [-v TOS]
[-r count] [-s count] [[-j host-list] ¦ [-k host-list]]
[-w timeout] target_name
```

The following list describes the switches available for use with PING:

- -t—Ping the specified host until stopped. To see statistics and to continue, type Ctrl+Break; to stop, type Ctrl+C.
- -a—Resolve addresses to hostnames.
- -n count—The number of echo requests to send.
- -l size—Send buffer size.
- -f—Set the Don't Fragment flag in the packet.
- -i TTL—Time to Live.
- -v TOS—Type of Service.
- -r count—Record route for count hops.
- -s count—Timestamp for count hops.
- -j host-list—Loose source route along host list.
- -k host-list—Strict source route along host list.
- -w timeout—Time in milliseconds to wait for each reply.

The ARP Command

The ARP command displays and modifies the IP-to-physical address translation tables used by Address Resolution Protocol (ARP), as shown here:

```
ARP -s inet_addr eth_addr [if_addr]
ARP -d inet_addr [if_addr]
ARP -a [inet_addr] [-N if_addr]
```

The following list describes the switches available for use with ARP:

- -a—Displays current ARP entries by referencing the current protocol data. If inet_addr is specified, the IP and physical addresses for only the specified computer are displayed. If more than one network interface uses ARP, entries for each ARP table are displayed.
- -g—Same as -a.
- inet_addr—Specifies an Internet address.
- -N if_addr—Displays ARP entries for the network interface specified by if_addr.
- -d—Deletes the host specified by inet_addr. You can use * as a wildcard with inet_addr to delete all hosts.
- -s—Adds the host and associates the Internet address inet_addr with the physical (MAC) address.
- eth_addr—Uses the physical (MAC) address and is given as six hexadecimal bytes separated by hyphens.
- if_addr—Specifies the Internet address of the interface that should have its address translation table modified. If if_addr is not entered, the first applicable interface is used.

For example, the following code adds a static entry:

```
> arp -s 157.55.85.212 00-aa-00-62-c6-09 ....
```

The following displays the ARP table:

```
> arp -a
```

The IPCONFIG Command

IPCONFIG is a command-line tool for getting basic IP configuration information, including the IP address, subnet mask, and default gateway. The IPCONFIG /all

switch produces a detailed configuration report for all interfaces on a system, including any configured remote access adapters, as shown here:

```
ipconfig [/? ¦ /all ¦ /renew [adapter] ¦ /release [adapter]
¦ /flushdns ¦ /displaydns ¦ /registerdns ¦ /showclassid adapter
¦ /setclassid adapter [classid] ]
```

The following list describes the switches available for use with `IPCONFIG`:

- `/all`—Display full configuration information.
- `/release`—Releases the IP address for the specified adapter.
- `/renew`—Renews the IP address for the specified adapter.
- `/flushdns`—Purges the DNS Resolver cache.
- `/registerdns`—Reregisters DNS names.
- `/displaydns`—Displays the contents of the DNS Resolver Cache.
- `/showclassid`—Displays all the DHCP class IDs allowed for adapter.
- `/setclassid`—Modifies the DHCP class ID.

The default is to display only the IP address, subnet mask, and default gateway for each adapter bound to TCP/IP. For `/release` and `/renew`, if no adapter name is specified, the IP address leases for all adapters bound to TCP/IP are released or renewed.

The `NBTSTAT` Command

NetBT Statistics (`Nbtstat.exe`) is a command-line tool that can be used to view and troubleshoot network NetBIOS over TCP/IP (NetBT) name resolution. It displays protocol statistics and current TCP/IP connections that are using NetBT.

Nbtstat resolves NetBIOS names to IP addresses by using several options for NetBIOS name resolution, including local cache lookup, WINS server query, broadcast, `LMHOSTS` and `HOSTS` file lookup, and DNS server query. It also displays protocol statistics and current TCP/IP connections using Nbtstat.

```
NBTSTAT [ [-a RemoteName] [-A IP address] [-c] [-n]
➥[-r] [-R] [-RR] [-s] [-S] [interval] ]
```

The following list describes the switches available for use with `NBTSTAT`:

- `-a` (adapter status)—Lists the remote machine's name table given its name.
- `-A` (adapter status)—Lists the remote machine's name table given its IP address.

- `-c` (cache)—Lists NBT's cache of remote (machine) names and their IP addresses.
- `-n` (names)—Lists local NetBIOS names.
- `-r` (resolved)—Lists names resolved by broadcast and via WINS.
- `-R` (Reload)—Purges and reloads the cache name table and reloads the `#PRE` tagged entries from the LMHOST file if any are present.
- `-S` (Sessions)—Lists the sessions table with the destination IP addresses.
- `-s` (sessions)—Lists the sessions table, converting destination IP addresses to computer NETBIOS names.
- `-RR`(ReleaseRefresh)—Sends Name Release packets to WINS and then starts Refresh.
- `RemoteName`—Remote host machine name.
- `IP address`—Dotted decimal representation of the IP address.
- *`interval`*—Redisplays selected statistics, pausing the number of seconds specified by *`interval`* between each display. Press Ctrl+C to stop redisplaying statistics.

The `NETSTAT` Command

`NETSTAT` (`Netstat.exe`) is a command-line tool that displays TCP/IP statistics and active connections to and from the local system. It can also display all connections and listening ports and has an option to display the number of bytes sent and received and any network packets dropped (if applicable).

```
NETSTAT [-a] [-e] [-n] [-o] [-s] [-p protocol] [-r] [interval]
```

The following list describes the switches available for use with `NETSTAT`:

- `-a`—Displays all connections and listening ports.
- `-e`—Displays ethernet statistics. Can be combined with the `-s` option.
- `-n`—Displays addresses and port numbers in numerical form.
- `-o`—Displays the owning process ID associated with each connection.
- `-p` *`protocol`*—Shows connections for the protocol specified by *`protocol`*, which can be `TCP`, `UDP`, `TCPv6`, or `UDPv6`. If used with the `-s` option to display per-protocol statistics, *`protocol`* can be any of the following: `IP`, `IPv6`, `ICMP`, `ICMPv6`, `TCP`, `TCPv6`, `UDP`, or `UDPv6`.

- `-r`—Displays the routing table.
- `-s`—Displays per-protocol statistics. By default, statistics are shown for IP, IPv6, ICMP, ICMPv6, TCP, TCPv6, UDP, and UDPv6; the `-p` option can be used to specify a subset of the default.
- *`interval`*—Redisplays selected statistics, pausing the number of seconds specified by *`interval`* between each display. Press Ctrl+C to stop redisplaying statistics. If this switch is omitted, `NETSTAT` prints the current configuration information once.

The `ROUTE` Command

The `ROUTE` command-line tool displays the current IP routing table for the local system, and it can be used to add or delete IP routes and to add persistent routes.

```
ROUTE [-f] [-p] [command] [destination] [MASK netmask]
➥[gateway] [METRIC metric] [IF interface]
```

The following list describes the switches available for use with `ROUTE`:

- `-f`—Clears the routing tables of all gateway entries. If it is used with one of the `ROUTE` commands (see the following list), the routing tables are cleared before running the command.
- `-p`—When used with the `ADD` command, it makes a route persistent across boots of the system. By default, routes are not preserved when the system is restarted. Ignored for all other commands, which always affect the appropriate persistent routes.

The following list describes the commands available for use with `ROUTE`:

- `PRINT`—Prints a route.
- `ADD`—Adds a route.
- `DELETE`—Deletes a route.
- `CHANGE`—Modifies an existing route.
- `destination`—Specifies the host.
- `MASK`—Specifies that the next parameter is the `netmask` value.
- `netmask`—Specifies a subnet mask value for this route entry. If not specified, it defaults to 255.255.255.255.

- `gateway`—Specifies the gateway.
- `interface`—Specifies the interface number for the specified route.
- `METRIC`—Specifies the metric—that is, the cost for the destination.

Names used for the `destination` command are looked up in the `NETWORKS` file on the local system. Names used for the `gateway` command are looked up in the `HOSTS` file on the local system. If the command is `PRINT` or `DELETE`, the destination or gateway can be a wildcard (`*`), or the gateway entry can be left blank. Invalid `MASK` entries, such as `(DEST & MASK) != DEST`, generate an error.

The `HOSTNAME` Command

`HOSTNAME` is a command-line tool for showing the local computer's hostname. It can be used for authentication purposes by the Remote Copy Protocol (RCP), Remote Shell (RSH), and Remote Execution (REXEC) tools.

The `TRACERT` Command

`TRACERT` is sometimes used to verify that IP addressing has been correctly configured on a client. It basically shows the route taken to reach a remote system.

```
tracert [-d] [-h maximum_hops] [-j host-list] [-w timeout] target_name
```

Here is a list of available switches for the `TRACERT` command:

- `-d`—Do not resolve addresses to hostnames.
- `-h` *`maximum_hops`*—Maximum number of hops to search for target.
- `-j host-list`—Loose source route along host list.
- `-w` *`timeout`*—Wait the number of milliseconds specified by *`timeout`* for each reply.

The `PATHPING` Command

Like `TRACERT`, `PATHPING` shows the route taken to reach a remote system, but `PATHPING` does so with more detail and offers more functionality.

```
pathping [-g host-list] [-h maximum_hops]
[-i address] [-n] [-p period] [-q num_queries]
[-w timeout] [-P] [-R] [-T] [-4] [-6] target_name
```

Here is a list of available switches for the `PATHPING` command:

- `-g` *`host-list`*—Loose source route along the host list.
- `-h` *`maximum_hops`*—Maximum number of hops to search for target.
- `-i` *`address`*—Use the specified source address.
- `-n`—Do not resolve addresses to hostnames.
- `-p` *`period`*—Wait the number of milliseconds specified by *`period`* between pings.
- `-q` *`num_queries`*—Number of queries per hop.
- `-w` *`timeout`*—Wait the number of milliseconds specified by *`timeout`* for each reply.
- `-P`—Test for RSVP PATH connectivity.
- `-R`—Test if each hop is RSVP aware.
- `-T`—Test connectivity to each hop with layer-2 priority tags.
- `-4`—Force using IPv4.
- `-6`—Force using IPv6.

The `FTP` Command

FTP is used to transfer files from system to system over TCP ports 20 and 21 (by default), but it can also help you diagnose problems on your TCP/IP network. By using Internet Explorer with FTP, users experience a Windows Explorer–type of GUI environment for the FTP file transfer by having features such as file and folder views, drag-and-drop, and copy-and-paste available.

The command-line `FTP` allows for more functionality. FTP is considered a connected session that uses TCP. `FTP` commands are as follows: `!`, `delete`, `literal`, `prompt`, `send ?`, `debug`, `ls`, `put`, `status append`, `dir`, `mdelete`, `pwd`, `trace ascii`, `disconnect`, `mdir`, `quit`, `type`, `bell`, `get`, `mget`, `quote`, `user`, `binary`, `glob`, `mkdir`, `recv`, `verbose`, `bye`, `hash`, `mls`, `remotehelp`, `cd`, `help`, `mput`, `rename`, `close`, `lcd`, `open`, and `rmdir`. Here is an example of the syntax:

```
FTP [-v] [-d] [-i] [-n] [-g] [-s:filename] [-a] [-w:windowsize] [-A] [host]
```

The following list explains the options you can use with the `FTP` command:

- `-v`—Suppresses display of remote server responses.
- `-n`—Suppresses auto-login at initial connection.

- `-i`—Turns off interactive prompting during multiple file transfers.
- `-d`—Enables debugging.
- `-g`—Disables filename globing.
- `-s:`*`filename`*—Specifies a text file containing FTP commands; the commands automatically run after FTP starts.
- `-a`—Use any local interface when binding a data connection.
- `-A`—Log in as anonymous.
- `-w:`*`buffersize`*—Overrides the default transfer buffer size of 4,096.
- *`host`*—Specifies the hostname or IP address of the remote host to connect to.

The `TFTP` Command

Trivial File Transfer Protocol allows for connectionless transfer of files to and from systems using UDP. Although TFTP is limited in functionality, there are still some command-line switches that can be used to tailor its performance:

```
TFTP [-i] host [GET ¦ PUT] source [destination]
```

Definitions for these switches are as follows:

- `-i`—Specifies binary image transfer mode (also called octet). In binary image mode, the file is moved literally byte by byte.
- *`host`*—Specifies the local or remote host by name.
- `GET`—Transfers the file destination on the remote host to the file source on the local host.
- `PUT`—Transfers the file source on the local host to the file destination on the remote host.
- *`source`*—Specifies the file to transfer.
- *`destination`*—Specifies where to transfer the file.

The `TELNET` Command

Telnet is a command-line terminal emulation program that enables an administrator to perform commands on a remote computer from a command window on a local system. Here is an example of the syntax:

```
telnet [-a] [-e char] [-f filename] [-l user] [-t term] [host] [port]
```

Definitions for TELNET switches are as follows:

- -a—Attempts automatic logon. Same as -l option, except it uses the currently logged on user's name.
- -e *char*—Escape character to enter Telnet client prompt.
- -f *filename*—Filename for client-side logging.
- -l *user*—Specifies the username to log in with on the remote system. Requires that the remote system support the TELNET ENVIRON option.
- -t *term*—Specifies terminal type. Supported term types are vt100, vt52, ansi, and vtnt only.
- *host*—Specifies the hostname or IP address of the remote computer to connect to.
- *port*—Specifies a port number or service name.

The RCP Command

Remote Copy Protocol (RCP) uses TCP to copy files to and from systems running the RCP service. It can be scripted in a batch file and does not require a password. The remote host must be running the Remote Shell Daemon (RSHD) service, and the user's username must be configured in the remote host's .rhosts file. Microsoft's implementation of TCP/IP includes the RCP client software but not RSHD services. RCP is one of the r-commands available on all UNIX systems.

```
RCP [-a ¦ -b] [-h] [-r] [host][.user:]source [host][.user:] path\destination
```

The following list explains the options you can use with the RCP command:

- -a—This is the default transfer mode that specifies ASCII transfer mode. This mode converts the end-of-line (EOL) characters to a carriage return for UNIX and a carriage return/line feed for personal computers.
- -b—Specifies binary image transfer mode.
- -h—Transfers hidden files.
- -r—Copies the contents of all subdirectories; destination must be a directory.
- *host*—Specifies the local or remote host. If the host is specified as an IP address *or* if the hostname contains dots, you must specify the user.

- `.user:`—Specifies a username to use instead of the current username.
- *source*—Specifies the files to copy.
- *path\destination*—Specifies the path to the logon directory on the remote host.

The RSH Command

Remote Shell (RSH) enables clients to run commands directly on remote hosts running the RSH service without having to log on to the remote host. Microsoft's implementation of TCP/IP includes the RSH client software but not the RSH service. If a user on a computer running in a Windows domain tries to use RSH to run a command on a remote UNIX server running RSH, the domain controller is required by the RSH client to resolve the user's username. `RSH` is one of the UNIX `r` commands that is available on all UNIX systems.

The REXEC Command

Remote Execution (REXEC) runs commands on remote hosts running the REXEC service and authenticates the username on the remote host before executing the specified command.

```
REXEC host [-l username] [-n] command
```

The following list defines options to use with the `REXEC` command:

- *host*—Specifies the remote host on which to run the command.
- `-l` *username*—Specifies the username on the remote host.
- `-n`—Redirects the input of `REXEC` to NULL.
- *command*—Specifies the command to run.

Planning a DHCP Strategy

As you learned earlier in the chapter, DHCP servers are required to implement dynamic IP addressing. The following section looks at some of the important topics you need to consider when planning a DHCP strategy.

DHCP Server Placement

Before you install the DHCP service, you need to identify where on the network the DHCP servers will be placed. Consider the following points when determining where you should place the DHCP servers:

- If there are slow network connections, consider placing DHCP servers locally. This will minimize traffic on the slow connections.
- Place DHCP servers as close as possible to DHCP clients to optimize response times.
- Consider using more than one DHCP server for fault tolerance.
- If there are multiple subnets, place DHCP servers on those subnets with the highest number of DHCP clients or those with frequent lease renewals.

DHCP Relay Agent

When a DHCP client request for an IP address hits a non-RFC-1542–compliant router (meaning that the DHCPDISCOVER broadcast message can't be forwarded off the subnet), it fails to receive a response because the DHCP server never receives the DHCPDISCOVER broadcast message, and the client system configures itself with an APIPA address.

If a DHCP Relay Agent is in use on the subnet, it receives the DHCPDISCOVER broadcast message and forwards (routes) the message off the subnet to the DHCP server. Subsequently, when the DHCP server responds with an address and the DHCP client selects the IP address, the client responds with a DHCPREQUEST broadcast message that includes the IP address of the server whose offer was accepted. Again, this DHCPREQUEST broadcast message does not get out of this subnet unless a DHCP Relay Agent in use on the subnet can receive the DHCPREQUEST broadcast message and forward (route) the message off the subnet to the DHCP server. Since the DHCPREQUEST is a broadcast, it will be blocked by routers that filter out broadcasts as a matter of their function being a layer 3 device. Thus, to allow a DHCP client to communicate with a DHCP server in a different subnet, a DHCP relay agent must be used.

80/20 Rule

DHCP servers should be located on at least one subnet on a LAN in a routed network. When the server needs to support clients on remote subnets separated by routers, the router should be RFC 1542–compliant, or a DHCP Relay Agent should be used to support forwarding DHCP traffic between the subnets.

When your network design allows for more than one DHCP server on the same subnet, the scope should be divided equally between them. If the DHCP servers are on separate subnets, the 80/20 rule should be deployed. With the 80/20 rule, the local DHCP server is configured with 80% of the IP addresses within the scope. The remote DHCP server is configured with the remaining 20%. This will allow client computers to continue to obtain valid IP addresses even if the local DHCP server stops responding. You need to make sure that the relay agent is enabled on your router for this configuration to work.

NOTE

The 80/20 rule for splitting addresses available in the scope can be tailored to meet your environment's addressing needs.

Centralized Versus Decentralized

Where DHCP servers are placed on the network depends on whether you require a centralized or decentralized DHCP infrastructure. With a centralized structure, the DHCP servers are placed in a central location. In a subnetted network, this requires DHCP Relay Agents or routers that are BOOTP enabled. One of the disadvantages to this model is that it can end up increasing network traffic on slow network connections.

With a decentralized structure, DHCP servers are placed on each subnet. This can reduce network traffic on slow network connections. However, it can increase costs because additional DHCP servers are needed.

You can also implement a combination of both infrastructures. After you define the physical characteristics of the network, you can identify those subnets that require a DHCP server. For example, a subnet with few DHCP clients might not require a DHCP server to be placed locally.

Client Reservations

In some instances, a workstation on the network requires a permanent IP address, but you still want that workstation to be a DHCP client. With a client reservation, the workstation can still be DHCP enabled, but the DHCP server always assigns the client the same IP address. The client reservation is based on the MAC address.

In terms of administration, the client network configuration settings remain the same, IP addressing remains centralized, and the clients can still be assigned optional parameters through the DHCP server.

NOTE

If multiple DHCP servers are configured with a range of IP addresses that cover the range of the reserved addresses, the client reservation must be duplicated on all DHCP servers. If not, the client might receive an incorrect IP address (one other than the preferred address reserved for the client on the first DHCP server).

DHCP Options

After a scope has been created, you can configure several DHCP options. The options can be configured at one of the following four levels:

- Server
- Scope
- Class
- Client

Options configured at the server level are applied to all DHCP clients, regardless of the subnet on which they reside. Any options that should be applied to all DHCP clients should be configured at this level. For example, to configure all clients on the network to use the same DNS server, you can configure the option at the server level. Keep in mind that when you are configuring scope options, any options configured at the scope or client levels override those configured at the server level. To configure server-level options, right-click the Server Options container listed under the DHCP server and select Set Predefined Options from the menu.

If you want to configure DHCP options so that they apply only to DHCP clients on a specific subnet, configure the options at the scope level. For example, the IP address of the default gateway for a subnet should be configured at the scope level. Configuring scope-level options can be done by right-clicking the Scope Options container and selecting Configure Options from the menu.

Finally, if you want to apply DHCP options to only a specific DHCP client, you can configure the options at the client level. You can configure options at this level only for clients that have a client reservation, meaning that they are DHCP clients but always lease the same IP address. Any option that you configure at this level overrides any configured at the server and scope levels. To configure a client-level option, right-click the client reservation and select Configure Options.

Windows Server 2003 also allows DHCP options to be applied to groups of users or workstations with similar needs. User-class options can be used to assign options to DHCP clients that have common needs for similar DHCP

options configurations. For example, a user class can be used to configure options for mobile users. Vendor-class options can be used to assign DHCP options on the basis of vendor information. For example, specific options can be assigned to clients running a specific version of Windows.

EXAM ALERT

Be sure you are familiar with the order in which scope options are applied. The order is server, scope, class, and then client.

Now that you're familiar with how DHCP options can be applied, let's take a look at the different DHCP options that can be assigned to clients. As previously mentioned, a DHCP server can assign parameters other than just an IP address and subnet mask to a DHCP client. A number of different options can be configured. To access the Server Options dialog box, highlight Server Options in the left pane of the DHCP management console and select Configure Options from the Action menu.

The following list provides a description of the commonly used DHCP options:

- *006 DNS Servers*—Specifies the IP address of the DNS servers available to clients on the network.
- *015 DNS Domain Name*—Specifies the DNS domain name used for client resolutions.
- *003 Router*—Specifies the IP address of the router or default gateway.
- *044 WINS/NBNS Servers*—Specifies the IP address of the WINS servers on the network available to clients.
- *046 WINS/NBT Node Type*—Specifies the name resolution type. The available options include 1 = B-node (broadcast), 2 = P-node (peer), 4 = M-node (mixed), and 8 = H-node (hybrid).
- *240 Classless Static Routes*—Specifies a list of static routes, including the destination network IP address, the subnet mask, and the router that is responsible for forwarding messages to that network.

Most of the options outlined in the preceding list can also be configured locally on the client. By doing so, any options configured on the DHCP server will be overwritten by those configured locally. If you are using DHCP, however, it would not make sense to configure the options locally as well, especially in terms of administrative overhead.

Securing DHCP

One of the ways that you can secure a DHCP implementation is to implement only Active Directory authorization. This would require all DHCP servers to be running Windows 2000 or later and be a member of an Active Directory domain. That way, when a DHCP server starts, it requests the server authorization list. If the DHCP server is not in the list of authorized servers, the DHCP service will fail to start. In order to authorize a DHCP server, your user account must be a member of the Enterprise Admins group.

Another way that you can secure your DHCP servers is to follow the principal of least privilege. Windows Server 2003 includes two built groups called the DHCP Users and the DHCP Administrators group. Members of the DHCP Users group have read-only access to the DHCP server while DHCP Administrators have full access. By placing users in the DHCP Users group, you can prevent unauthorized changes from being made to the server.

Optimizing DHCP

Each DHCP scope is configured with a lease duration. This specifies how long a DHCP client can use an IP address before it must be renewed by a DHCP server. By default, this value is set to eight days. However, you might want to change this depending on the number of IP addresses available as compared to the number of DHCP clients.

The lease duration can be customized to meet the requirements of your network. If the number of IP addresses exceeds the number of DHCP clients on the network, you can configure a longer lease duration. However, if the number of IP addresses available in the scope is comparable to the number of DHCP-enabled clients, you should configure a shorter lease duration. Also, if your network consists of a number of mobile users who move between subnets, consider creating a shorter lease time. By shortening the lease duration, you might also see a slight increase in network traffic because IP addresses are renewed at a more frequent interval.

Plan a Host Name Resolution Strategy

The Windows Server 2003 DNS Service offers a number of new features and enhancements, and many of the Windows 2000 features have been carried over as well. For example, you can configure *conditional forwarders* on your Windows Server 2003 DNS server to forward all DNS queries for a specific domain name to an IP address of a specific DNS server or servers. Conditional forwarders can be used in both intranet and Internet queries.

You can also use forward-only servers when you need to manage the DNS traffic between your network and the Internet. To do this, you configure the firewall your network uses so that only one DNS server is allowed to communicate with the Internet. This requires you to configure the other DNS servers in your enterprise to forward queries they cannot resolve locally to the Internet-enabled DNS server in the DMZ so that the query can be resolved.

The Windows Server 2003 DNS service provides basic support of the DNS Security Extensions (DNSSEC) protocol defined in RFC 2535, which allows DNS servers to perform as secondary DNS servers for existing DNSSEC–compliant, secure zones. Windows Server 2003 DNS supports storing and loading DNSSEC-specific resource records (RRs).

Microsoft DNS on Windows Server 2003 is compliant with most of the RFC specifications used to define the DNS protocol and allows deploying Active Directory under other DNS implementations. The Windows Server 2003 DNS service supports the following features:

- IETF Internet-Draft "A DNS RR for specifying the location of services (DNS SRV)" (SRV records)
- Dynamic updates
- Secure dynamic update based on the General Security Service Transaction Signature (GSS-TSIG) algorithm
- WINS and WINS R (reverse) records
- Fast zone transfer
- Incremental zone transfer
- Support for UTF (eight-character encoding)

NOTE

Berkeley Internet Name Domain (BIND) DNS servers regard Active Directory–integrated zones as Standard Primary DNS zones. Active Directory–integrated zones can replicate DNS updates to other Active Directory–integrated zones or to Standard Secondary DNS zones. Because Active Directory–integrated zones can replicate DNS data to Standard Secondary DNS zones, you can use Active Directory–integrated zones with BIND servers hosting Standard Secondary DNS zones.

DNS is the primary naming convention for Windows 2000 and 2003 domains; it provides name resolution for client systems by translating computer names to IP addresses so that computers can locate each other. DNS domains and Active Directory domains can share a common naming structure; in many cases, they

are identical, but they can also be completely different. For example, Server1.gunderville.com is a valid Windows domain name and could be the internal name for a host. If that same server were available to the Internet for access, it could also use that naming convention if it was available. The best analogy for correlating DNS names with IP addresses is using the phone book to look up someone's name (the DNS name, in other words) to find his or her area code and phone number (that is, the IP address).

Understanding Name Resolution

There are two types of DNS lookup queries: forward and reverse. A *forward lookup query* resolves a DNS name to an IP address and is the most common DNS query. When you perform a forward lookup, such as entering http://www.zandri.net into a browser, your client looks up the website's IP address with the assistance of a DNS server behind the scenes. A *reverse lookup query* resolves an IP address to a name. A DNS name server can resolve a query only for a zone for which it has authority. When DNS servers receive a name resolution request, they attempt to locate the requested information in their own cache and local database.

DNS servers cache all external name resolution data for a specific interval called Time to Live (TTL). The default TTL for DNS resolution is one hour; for WINS lookups via DNS, the default is 15 minutes.

DNS administrators can adjust the default settings for the DNS cache by going to the zone's Properties dialog box and selecting the Start of Authority (SOA) tab. To change the 15-minute default setting for WINS, select the WINS tab and click the Advanced Settings button.

Configuring a shorter TTL interval ensures that DNS and WINS information is more up to date, but it also increases the load on your DNS server, your WINS server, and your network. You can increase the interval when network resources are a higher premium than "freshness" of name resolution.

Two types of queries can be performed in DNS: iterative and recursive. An *iterative query* happens when a client makes a DNS resolution query to a DNS server, and the server returns the best answer it can provide based on its local cache or stored zone data. If the server resolving the iterative query does not have an exact match for the name request, it provides a pointer to an authoritative server in another level of the domain namespace. The client system then queries that server, and continues this process until it locates a server that is authoritative for the requested name or until an error is returned, such as "name not found," or a timeout condition is met.

A *recursive query* happened when a client makes a DNS resolution query to a DNS server, and the server assumes the full workload and responsibility for providing a complete answer to the query. If the server cannot resolve the query from its own database, it performs separate iterative queries to other servers (on behalf of the client) to assist in resolving the recursive query. The server continues this process until it locates a server that is authoritative for the requested name or until an error is returned, such as "name not found," or a timeout condition is met.

In most cases, client computers send recursive queries to DNS servers, and usually the DNS server is set up to make iterative queries to supply an answer to the client. In the following query example, a client computer makes a request to a DNS server to resolve the web address http://www.zandri.net:

1. The client computer generates a request for the IP address of www.zandri.net by sending a recursive query to the DNS server it is configured to use in its network configuration. (Call this server LOCALCFG.)
2. The LOCALCFG DNS server looks in its local database for an answer. If it finds that answer locally, it is returned to the client. Usually this happens only if the server is authoritative for the DNS zone in question or if it is hosting secondary zone files for other DNS zones. If the client is client1.zandri.net and is looking for www.zandri.net from DNS server localcfg.zandri.net, localcfg.zandri.net would likely have the required information in its own database.

   For the purposes of the remainder of this description, assume that the client looking for http://www.zandri.net is not going to find DNS resolution on the local DNS server and that LOCALCFG is not a member of the zandri.net domain and does not host secondary zone files for the zandri.net domain.

NOTE

The answer to this DNS query might be in the local cache if the DNS server recently looked up this resolution request for another client. In a large enterprise with many DNS clients connecting to the Internet, it would not be uncommon for many of the largest Internet sites to be almost constantly present in the local DNS server's cache because clients throughout the enterprise commonly query for the name resolution of those servers.

3. If LOCALCFG is unable to locate an entry for www.zandri.net in its own database, it sends an iterative query to a DNS server that is authoritative for the root of the local domain. (Call this server LOCALROOT.)

4. If the LOCALROOT DNS server, which is authoritative for the root domain, has the answer in its local database, it sends a response to LOCALCFG. If the LOCALROOT DNS server is unable to locate an entry for www.zandri.net in its database, it sends a reply to the LOCALCFG DNS server with the IP address of a DNS server that is authoritative for the .net domain. (Call these servers DNSNET*x*; *x* would be the numerical designation of a particular server.)

NOTE

If the address ends in `.com`, IP addresses of DNS servers that are authoritative for the `.com` domain are sent. For addresses ending in `.org`, IP addresses of DNS servers that are authoritative for the `.org` domain are sent, and so on.

5. The LOCALCFG DNS server that received the client's recursive query sends an iterative query to the DNSNET server.
6. If the DNSNET server has an entry for www.zandri.net in its local cache, it returns the answer to the LOCALCFG DNS server. If DNSNET is unable to locate an entry for www.zandri.net in its database, it sends a reply to the LOCALCFG DNS server with the IP addresses of DNS servers that are authoritative for the zandri.net domain. (Call this server ZANDRIDNS.)
7. The LOCALCFG DNS server sends an iterative query to the ZANDRIDNS server.
8. The ZANDRIDNS server locates an entry for www.zandri.net in its database and sends a reply to the LOCALCFG DNS server with the IP address of www.zandri.net.
9. The LOCALCFG DNS server sends a reply to the client computer with the IP address of www.zandri.net, which allows that client's web browser to display the web page on www.zandri.net.

When DNS clients make a request for a reverse DNS lookup, they are effectively making a request to resolve a hostname of a known IP address. In the standard DNS namespace, there is no connection between hostnames and IP addresses, and only a thorough search of all domains allows for reverse resolution.

Often, DNS servers are called on to resolve the same query multiple times within a short time. As an example, if a number of AOL users (arguably the largest ISP in the world) get an email reporting that new articles have been posted to www.2000trainers.com and immediately go to this site to read the new articles, the AOL DNS servers must continually recall the resolved address many times within a short period and use their local cache to do so.

DNS servers cache resolved addresses for a specific duration, specified as the TTL in the returned data. The DNS server administrator of the zone containing the data decides on the TTL setting. Therefore, the administrator of the 2000trainers.com domain and the DNS servers for that domain sets the TTL value. This value tells the resolving AOL DNS servers how long to hold that information in their cache. The lower the TTL, the "fresher" the resolution data is on the resolving AOL DNS servers.

After a DNS server caches data, it decreases the TTL from its original value so that it knows when to flush data from its cache. If another resolution query comes to the DNS server for the URL, the cached data is used and the TTL is reset (in most cases) to the original TTL. (The only time the TTL value isn't reset to its original value is if the administrator sets a different TTL.)

NOTE

In some instances, it is practical and perhaps necessary to disable recursion on certain DNS servers. If your enterprise has a number of internal-only DNS servers, it is not necessary for those servers to continue attempting to resolve yahoo.com for your clients, for example.

By setting certain DNS servers to not perform recursive lookups, in effect external lookup resolutions would quickly fail (as they would not have local data for yahoo.com). This allows client systems to fail over to another DNS server they have been configured to use to resolve the external name. For this process to be effective, the other DNS server would have to be one that handles external lookups.

The in-addr.arpa domain was created to avoid query overloads on DNS servers. System names in the in-addr.arpa domain are listed by their respective IP addresses. Because IP addresses are designed so that they become more significant as you move from left to right in the address, and domain names get less significant from left to right, IP addresses in the in-addr.arpa domain are listed in reverse order.

Pointer (PTR) records are added to IP addresses and corresponding hostnames. To perform a successful reverse lookup of an IP address, such as 121.41.113.10, the DNS server performing the query looks for a PTR record for 10.113.41.121.inaddr.arpa, which has the hostname and IP address 121.41.113.10.

Planning DNS Servers

Before you can effectively implement DNS, you need to take some time to plan the infrastructure. This includes determining where your DNS servers are placed, how many DNS servers are required, what type of DNS server role best meets your requirements, and so on. The following section looks at some of the important topics that need to be considered when implementing DNS.

DNS Server Placement

There are two points that need to be considered where considering DNS Server placement. First, how many DNS servers do you require and second, where on the network will they be located.

DNS servers should be placed in a location where they are most accessible to DNS clients. Ideally, a DNS server should be placed on each subnet. When it comes to the number of DNS servers, you should have a minimum of two DNS servers for your zone. However, this will depend on the level of fault tolerance that you require.

DNS Server Roles

You can configure a DNS server in one of three possible roles: caching-only DNS server, primary DNS server, or secondary DNS server. The role the server plays depends on the configuration of *zone files* and how they are maintained. The zone files contain configuration information for the zone as well as the resource records.

Caching-only DNS servers perform name resolution on behalf of clients and then cache the resulting name resolutions. They are not configured to be authoritative for a DNS zone and do not store Standard Primary or Standard Secondary zones locally. Their local cache holds the most frequently requested names and associated IP addresses and are available for subsequent client queries. This helps reduce traffic across a WAN because a caching-only server attempts to locate information in its cache to resolve local client requests and queries across the WAN for name resolution only when the information needed to complete the name resolution request is not available.

Also, because a caching-only server does not maintain any zone files that need to be updated via replication, it does not generate any zone transfer traffic.

All DNS servers maintain a `cache.dns` file that contains a list of all Internet root servers. Any time a DNS server resolves a hostname to an IP address, the information is added to the cache file. The next time a DNS client needs to resolve that hostname, the information can be retrieved from the cache instead of the Internet.

A primary DNS server hosts the working (writable) copy of a zone file. If you need to make changes to the zone file, it must be done from the server that is designated as the primary server for that zone. For those of you who are familiar with Windows NT 4.0, this is similar to how the primary domain controller (PDC) maintains the working copy of the directory database. After a server is configured as a primary DNS server for a zone, it is said to be authoritative for that domain. In addition, a single DNS server can be the primary DNS server for multiple *zones*.

A secondary server gets all its zone information from a master DNS server. The secondary DNS server hosts a read-only copy of the zone file, which it gets from the primary server or another secondary DNS server. Through a process known as a *zone transfer*, the master DNS server sends a copy of the zone file to the secondary server. The server responsible for the transfer of information to a secondary DNS server is known as a Master server.

NOTE

Pre–Windows 2000 implementations of DNS supported only full transfers or AXFR, in which an update to the zone file resulted in the entire zone database being transferred to the secondary servers. Windows Server 2003 (as well as Windows 2000 DNS) supports incremental zone transfers or IXFR, so the secondary servers can synchronize their zone files by pulling only the changes. This results in less network traffic.

For example, if Server2 is configured as a secondary server for bayside.net, Server2 would get all of its zone information from Server1, the primary DNS server for the zone. Any changes that need to be made to the zone file would have to be done on Server1. The changes would then be copied to Server2. As already mentioned, a DNS server can be both a primary and a secondary server at the same time. Using this example, Server2 could also be configured as the primary server for riverside.net, and, to provide fault tolerance for the zone file, Server1 could be configured as a secondary server for this zone.

Secondary DNS servers provide the following benefits:

- *Fault tolerance*—Because the secondary server has a copy of the zone file, name resolution can continue if the primary DNS server becomes unavailable.
- *Reduction in name-resolution traffic*—Secondary servers can be placed in remote locations with a large number of users. Clients can then resolve hostnames locally instead of having to contact a primary DNS using a WAN link.
- *Load balancing*—Name-resolution services for a zone can be provided by the secondary server as well, thereby reducing the load placed on the primary DNS server.

DNS Zone Overview

A *DNS zone* is a contiguous portion of the domain namespace for which a DNS server has authority to resolve DNS queries. DNS namespaces are almost

always divided into zones, which store name information about one or more DNS domains or portions of a DNS domain.

In the Windows Server 2003 Active Directory domain structure, there are three different zone types: Standard Primary zones, Standard Secondary zones, and stub zones.

Standard Primary Zones

Standard Primary zone files contain a read/write version of the zone file that is stored in a standard text file. Any changes to the zone are recorded only in that file. Any other copies of that zone are Standard Secondary zone copies and are read-only.

Standard Secondary Zones

A *Standard Secondary zone* file contains a read-only version of a Standard Primary zone file stored in a standard text file. Any changes to the zone are performed on the Primary zone file and replicated to the Secondary zone file. You create a Standard Secondary zone to make a copy of an existing Standard Primary zone and its zone file, which allows the DNS name resolution workload to be distributed among multiple DNS servers, thus providing a certain level of fault tolerance and load balancing. In addition, for remote sites, a Standard Secondary zone on that local site keeps local user systems from consuming unnecessary bandwidth by going over the WAN to query the DNS server. In most cases, local user systems are configured with the local DNS server that hosts the Standard Secondary zone as their primary DNS server, and perhaps another DNS server (at the main site) as a secondary DNS server.

Stub Zones

Stub zones allow a DNS server that is authoritative for the parent zone to be kept aware of other authoritative DNS servers for its child DNS zones in an effort to maintain DNS name resolution parity. Configuring stub zones enables administrators to distribute a list of the authoritative DNS servers for a zone without using Standard Secondary zones. Stub zones do not serve the same purpose as Standard Secondary zones and should not be used explicitly for redundancy and load balancing.

Active Directory Integrated Zones

Active Directory Integrated zones store DNS zone information in the Active Directory database rather than a text file. The Active Directory Integrated option is not available in the Change Zone Type dialog box unless the DNS server is also a domain controller. If the DNS server is not configured as a domain controller in your environment, the option to create the zone as Active Directory Integrated is grayed out during the creation of the zone (see Figure 3.3).

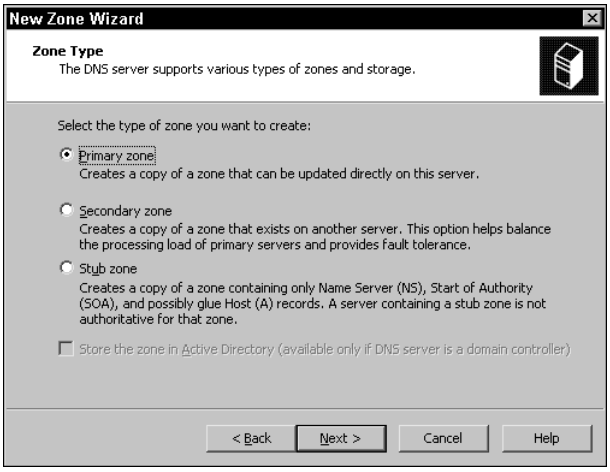

FIGURE 3.3 The option to create the DNS zone as Active Directory Integrated is grayed out because the server is not installed in the role of a domain controller.

Zone Data Replication

When replication is considered for Standard zone types, Primary zones are the only read/write copy of the DNS zone database, and only one Standard Primary zone exists. You can replicate the zone to any number of Standard Secondary DNS zones as needed, and these zones are read-only copies of the zone information. The only place where updates to the DNS zone database can occur is on the Standard Primary copy of the zone. The changes recorded there are replicated incrementally by incremental zone transfer (IXFR) or by transferring the entire zone database via full zone transfer (AXFR).

Although a Standard Primary DNS server with multiple Standard Secondary servers is fault tolerant and load-balanced to a degree, the loss of the Standard Primary DNS server prevents updates to the DNS database until an administrator can intervene. The administrator needs to repair the downed Standard Primary DNS server or promote a Standard Secondary DNS server to a Standard Primary DNS server. Name resolution still occurs for the first 24 hours with only the Standard Secondary servers functioning then all the records on the secondary servers will expire.

DNS zone transfers occur on a preconfigured basis, which is set to 15 minutes by default. This preconfigured basis is called a *refresh interval*, and its setting is found in the Start of Authority (SOA) tab of a zone's Properties dialog box. When the refresh interval expires, DNS servers request a copy of the current SOA record for their zone.

The DNS server then compares the serial number of the source DNS server's current SOA record with the serial number in its own local SOA record. If the responding DNS serial number is higher, the DNS server requests a zone transfer from the Primary DNS server. If a network issue or some other error prevents

the zone transfer, the requesting DNS server retries the request for a zone transfer at a preconfigured interval, which is called a *retry interval* and is set to 10 minutes by default.

If the failure is persistent and a zone transfer cannot be completed, the requesting DNS server quits making requests for a zone transfer after the *expire interval* (default setting is one day) has been exceeded. Active Directory–integrated DNS zones store zone information in Active Directory and are multimaster in nature, which enables administrators to configure systems so that updates to the DNS database can occur on any domain controller hosting the DNS zone. This type of setup is fully load balanced and fault tolerant; the loss of any single DNS server does not affect DNS name resolution or updates to the DNS database because any DNS server can receive updates.

Because Active Directory–integrated zones store zone information in Active Directory, the DNS information is replicated along with other Active Directory data. This configuration also enables administrators to establish permissions using Access Control Lists (ACLs) to allow only authenticated systems, groups, or users to make updates to the DNS zone.

NOTE

ACL permissions to update the DNS zone are made in the DNS zone container within Active Directory and can be assigned for the entire DNS zone or for individual resource records within the zone. They can also be assigned to systems, groups, or users.

Active Directory–integrated zones can be configured for the following types of replication (see Figure 3.4):

- Replication can be configured so that all DNS servers in the Active Directory forest can replicate DNS zone data to all DNS servers running on domain controllers within the Active Directory forest. (For the most part, this is the Active Directory–integrated DNS zone replication in Windows 2000 Server.)
- Replication can be configured so that all DNS servers in the Active Directory domain can replicate DNS zone data to all DNS servers running on domain controllers in the local domain. This is the default Active Directory–integrated DNS zone replication setup in Windows Server 2003.
- Replication can be configured so that all domain controllers in the Active Directory domain can replicate DNS zone data to all domain controllers in the Active Directory domain. You can use this setting if you need Windows 2000 DNS servers to load a specific Active Directory zone.

- Replication can be configured so that all domain controllers in a specified application directory partition can replicate DNS zone data according to the replication scope of the specified application directory partition. For a zone to be stored in the specified application directory partition, the DNS server hosting the zone must be enlisted in the specified application directory partition.

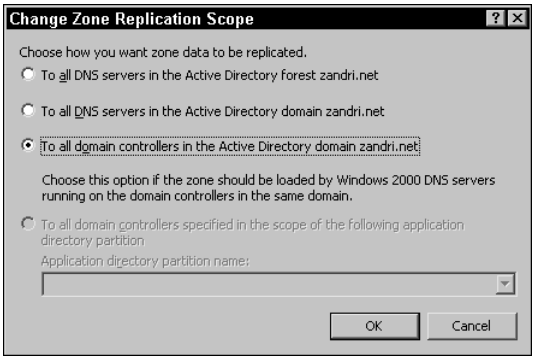

FIGURE 3.4 There are three options for changing the scope of DNS replication in Active Directory.

NOTE

Active Directory–integrated DNS zone data replicated to application directory partitions is not replicated to the forest's Global Catalog. The domain controller containing the Global Catalog can also host application directory partitions, but it does not replicate this data to its Global Catalog.

Active Directory–integrated DNS zone data replicated to domain partitions is replicated to all domain controllers in its Active Directory domain, and a portion of this data is stored in the Global Catalog. This setting is used to support Windows 2000.

TIP

The default configuration of the Windows Server 2003 DNS service allows zone transfers only to servers listed in a zone's Name Server (NS) resource records. Administrators can increase the security of the zone transfer process by changing the setting to allow zone transfers to the specified IP addresses of those servers only.

If DNS zone transfers must take place over a public network, such as the Internet, you can protect replication data by using VPNs and IPSec encryption. Using Active Directory–integrated zones exclusively for replicating DNS zone data adds another layer of security to zone transfers.

DNS zone files contain the name resolution data for a zone and include resource records with database entries containing various attributes of network systems. The most common resource records are discussed in the following sections.

(A) Records

Sometimes called host records or address records, *(A) records* contain the name–to–IP address mapping information used to map DNS domain names to a host IP address on the network. The following are examples of host or (A) records:

```
server1         IN A 121.41.113.10
localhost       IN A 127.0.0.1
```

Alias Records

Normally referred to as CNAME (canonical name) records, *Alias records* allow you to provide additional names to a server that already has a name in an (A) resource record. This is how a web server with a name of Server1 in a domain of zandri.net "becomes" www.zandri.net, as far as DNS resolution is concerned: There is an Alias record referencing www.zandri.net to Server1.zandri.net. Here are some examples of using CNAME records:

```
www             CNAME Server1
ftp             CNAME Server1
```

MX (Mail Exchanger) Records

MX (Mail Exchanger) records specify which server email should be delivered to in a domain. When you have a mail server named Mailbox.zandri.net and you want all mail for all_users@zandri.net to be delivered to this mail server (named Mailbox in this example), the MX resource record must exist in the zone for Zandri.net and point to Mailbox.

NS (Name Server) Records

NS (Name Server) records designate DNS domain names for the servers that are authoritative for a DNS zone and can list additional name servers within the record. The following is an example of an NS record:

```
@ IN NS server2.zandri.net
```

The at symbol (`@`) in a database file indicates "this server" and the `IN` indicates an Internet record.

PTR (Pointer) Records

PTR (Pointer) records are used for reverse lookup queries that resolve IP addresses to names. Reverse lookup zones are created in the in-addr.arpa domain to designate a reverse mapping of a host IP address to a host DNS domain name.

As mentioned previously, to perform a successful reverse lookup of an IP address, such as 121.41.113.10, the DNS server performing the query looks for a PTR record for 10.113.41.121.in-addr.arpa, which has the hostname and IP address 121.41.113.10. The PTR record for it looks like this:

```
10.113.41.121.in-addr.arpa. IN PTR Server1.Zandri.net.
```

Reverse lookup zones are not a requirement; they are an optional configuration for your DNS server in most cases. In certain situations, a reverse lookup zone might be required to verify the location of connecting clients.

EXAM ALERT

The Nslookup command will not work if you do not have reverse lookup records.

SOA (Start of Authority) Records

SOA (Start of Authority) records indicate the starting point of authority for a DNS zone on a specific DNS server. The SOA record is the first resource record created when you add a new zone. The following is an example of an SOA record:

```
@ IN SOA server1.zandri.net. (
                                1        ; serial number
                                7200   ; refresh [2h]
                                900     ; retry [15m]
                                86400 ; expire [1d]
                                7200 ) ; min TTL [2h]
```

The at symbol (@) in a database file indicates "this server." `IN` indicates an Internet record. Any hostname not terminated with a period has the root domain appended. The @ symbol is replaced by a period in the administrator's email address. Parentheses must enclose line breaks that span more than one line. The `7200 ; refresh [2h]` shows a refresh interval of two hours, `900 ; retry [15m]` shows a retry interval of 15 minutes, `86400 ; expire [1d]` shows an expire interval of one day, and `7200 ; min TTL [2h]` shows a minimum TTL of two hours. Everything after a semicolon (;) is a comment, which is why line breaks are necessary.

SRV (Service) Records

Sometimes referred to as Service Location records, *SRV (Service) records* contain registered services within the zone so that clients can locate these services by using DNS. SRV records are mainly used to identify services in Active Directory.

The `CACHE.DNS` file contains the records of the root DNS servers. The cache file is basically the same on all name servers and must be present for a DNS server to handle a query outside its zone. The file provided by default with the Windows 2003 DNS server has current records for all the root servers on the Internet. This file is stored in the `%SystemRoot%\System32\Dns` folder when you install DNS on your Windows Server 2003 server.

If you are running DNS for internal use, not for connections for forwarding to the Internet, the `CACHE.DNS` file should be replaced to contain the name server's authoritative domains for the root of the private network. In certain situations, you replace the `CACHE.DNS` file in the `%SystemRoot%\System32\Dns` folder, and it does not update the root hints listed in the DNS Manager. This can happen because the DNS server is a domain controller configured to load zone data on startup from Active Directory and the Registry. If the root hints specified in Active Directory have been deleted, modified, incorrectly entered, or damaged, this behavior occurs. There is more information on `CACHE.DNS` in "Need To Know More?" at the end of this chapter.

Root hints consist of a list of resource records that the DNS service can use to locate other DNS servers that are authoritative for the root of the DNS domain namespace tree in your enterprise. For DNS servers that have been deployed to resolve names external to your environment, root hints host addresses for the Internet root servers. When a new server is added or removed from your DNS structure, administrators might need to update the root hints list.

For a current copy of the root hints file for Internet root servers, you can download a copy of the `named.root` file from ftp://ftp.internic.net/domain/. Figure 3.5 shows the DNS files you can download from InterNIC for DNS use.

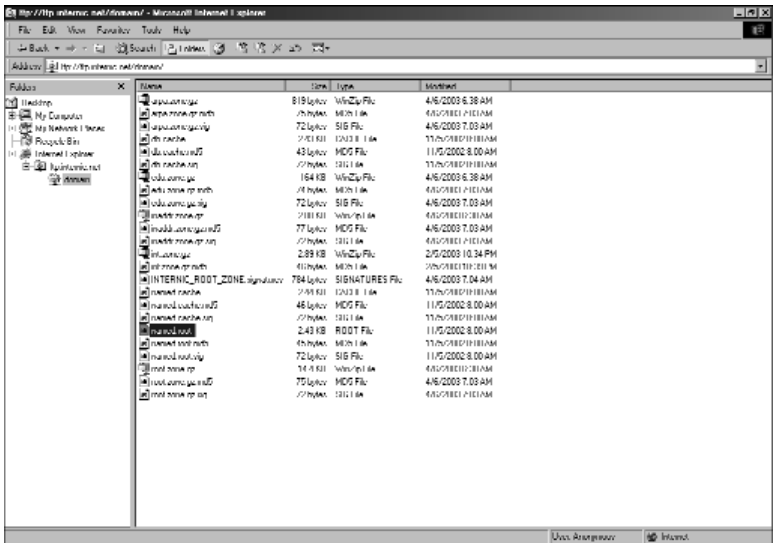

FIGURE 3.5 You can download DNS files from InterNIC to use with DNS.

Choosing a DNS Namespace

You must take a number of considerations into account when you are planning your DNS namespace design. When setting up DNS in your environment, you should register a unique parent (second-level) DNS domain name that can be used for hosting your organization on the Internet, such as "zandri" in zandri.net or "2000trainers" in 2000trainers.com. Even if you are not currently connected to the Internet via your proposed domain name or are not considering it in the near future, you should still make an effort to register your DNS domain name because you can't accurately predict what course your company might take in the future. If your company does decide to establish an Internet presence, you might find that the name you are currently using is registered to someone else. In that case, not only will you be unable to use it, but you will have to redesign your entire DNS deployment.

Also, you need to consider to which top-level domain name you will tie your second-level name. Table 3.3 shows the original top-level domain names currently available on the Internet. Keep in mind that 2 letter country codes and many new top level domain names are not available.

TABLE 3.3 Top-Level Domain Names

Top-Level Name	Description	Traditionally Used By
arpa	Run by Advanced Research Project Agency (ARPA). Used to register reverse mapping of IPv4 addresses assigned to DNS domain names for computers that use those addresses on the Internet.	The in-addr.arpa domain
com	For business and commercial use.	Businesses and corporations
edu	For educational use.	Public and private schools, colleges, and universities
gov	For use by government institutions.	Local, state, and federal government agencies
int	Reserved for international use. Currently planned for use in RFC 1886 to register reverse mapping of IPv6 addresses assigned by IANA to DNS domain names in the ip6.int domain for computers that use those addresses on the Internet.	The ip6.int domain

TABLE 3.3 *Continued*

Top-Level Name	Description	Traditionally Used By
mil	For use by military agencies.	Department of Defense (DoD), U.S. Navy, U.S. Army, U.S. Air Force, and other military agencies
net	For use by organizations that provide large-scale Internet or telephony-based service.	Large-scale Internet and telephone service providers
org	For use by noncommercial, nonprofit organizations.	Noncommercial, nonprofit organizations and charitable institutions

After your parent domain name is in place, you can configure other child names as needed. For example, forums.2000trainers.com can be created as a child of 2000trainers.com for resources the forums section of 2000trainers uses, and additional child domains could be established as needed. The domain might be broken down by geographical region, in which us.forums.2000trainers.com is one child domain and canada.forums.2000trainers.com is another.

Active Directory domains often correspond directly with DNS names. When choosing DNS names to use for Active Directory, starting with the registered DNS domain name suffix your organization has reserved for use on the Internet is best. Effectively, it means that internal and external namespaces are the same. For example, Zandri.net is what you can reach from the Internet *and* the intranet. The main benefit is that users access one domain name when they need to find resources, regardless of where those resources exist.

This common naming scheme causes a certain amount of additional administration to make certain that appropriate DNS and SRV records are stored on internal and external DNS servers. You could also use a delegated namespace internally, so that 2000trainers.com is reachable from the Internet and the Active Directory namespace is set up with something along the lines of corp.2000trainers.com.

This configuration offers better security because users and computers outside the organization cannot access the private DNS namespace from standard Internet connections with the proper segmentation of DNS zones. There is also less administrative overhead for DNS and SRV records because the zones for both namespaces are independent of each other. This setup could cause some confusion for internal users, however, if they need to access certain network resources from the external domain.

Optionally, the entire internal Active Directory namespace can be totally separate from the external namespace, so that zandri.net is used from the Internet

and mcmcse.com is used internally. This configuration also offers better security because users and computers outside the organization cannot access the private DNS namespace from standard Internet connections with the proper segmentation of DNS zones. There is also less administrative overhead for DNS and SRV records because the zones for both namespaces are independent of each other. This setup could also cause confusion for internal users, however, if they need to access certain network resources from the external domain.

Zone Delegation

Delegation is the process of designating a portion of the DNS namespace for another zone. It gives administrators a way of dividing a namespace among multiple zones. For example, an administrator might place the bayside.net domain in one zone and place the sales.bayside.net subdomain in another delegated zone. The bayside.net zone would contain all the records for the sales subdomain if it is not delegated. Through delegating, the bayside.net zone contains only information for bayside.net, as well as records to the authoritative name servers for the sales.bayside.net zone. The host entries for any machines in sales.bayside.net are contained only on the delegated server.

In any case, when deciding whether to delegate, keep the following points in mind:

- Zone delegation allows you to delegate management of part of the DNS namespace to other departments or locations.
- Zone delegation allows you to distribute a large DNS database across multiple servers for load balancing, faster name resolution, and increased performance.
- Zone delegation allows you to extend the namespace for business expansion, that is, it is scalable with business needs.

NOTE

To facilitate the delegation of zones, you need the appropriate delegation records that point to authoritative name servers for the new zone(s).

Optimizing DNS

Administrators can optimize DNS servers in the enterprise by disabling local subnet prioritization and round-robin rotation. *Local subnet prioritization* is used so that clients on the same subnet as the available DNS server, based on IP address location, have priority over other DNS clients. *Round-robin rotation* of

available DNS servers provides network load balancing. Disabling either or both settings reduces and balances client response time, for the most part, across the entire enterprise. Both settings are enabled by default.

Other DNS configurations can be modified from their default settings in an effort to tailor preferences to a locale's specific needs. To adjust these settings, right-click the DNS server in the DNS MMC, choose Properties, and select the Advanced tab of the Properties dialog box. Table 3.4 shows the properties that can be configured and the default settings.

TABLE 3.4 Default DNS Settings in Windows 2003 Server

Property	Default Settings
Disable recursion	Off
BIND secondaries	On
Fail on load if bad zone data	Off
Enable round robin	On
Enable netmask ordering	On
Secure cache against pollution	On
Name checking	Multibyte (UTF8)
Load zone data on startup	From Active Directory and Registry
Enable automatic scavenging of stale records	Off

Client Dynamic Updates and DHCP

You can configure DHCP servers in your enterprise to dynamically update DNS when the DHCP server assigns a DHCP client computer IP information, or you can allow clients to dynamically update DNS.

DNS clients running Windows 2000, Windows XP, and Windows Server 2003 operating systems can dynamically update DNS on startup. When DNS clients are allowed to update DNS, they connect to the DNS server on startup and automatically register the appropriate client information, such as the system IP address and the fully qualified domain name (FQDN), with the DNS server, regardless of whether their IP addresses are entered manually or assigned via DHCP.

To have clients dynamically update DNS, in the Network dialog box, select the client's active network connection, and choose Properties. Select Internet Protocol (TCP/IP), and click the Properties button. In the General tab, click the Advanced button to open the Advanced TCP/IP Settings dialog box, and then select the DNS tab (see Figure 3.6).

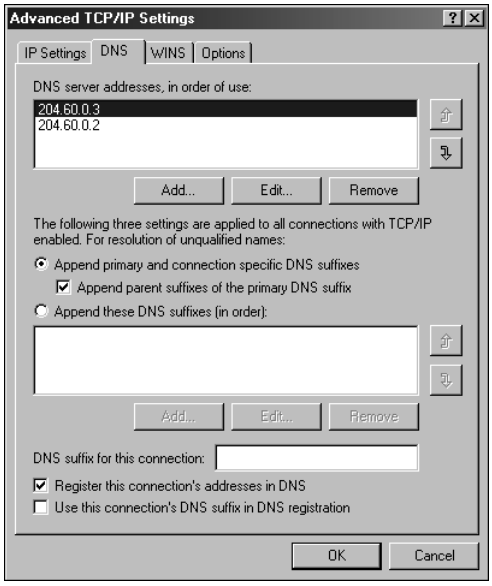

FIGURE 3.6 The Register this connection's addresses in DNS check box is enabled at the bottom of the Advanced TCP/IP Settings dialog box.

DNS dynamic update is enabled by default. You can configure additional dynamic update settings with the DHCP console. To do this, right-click the DHCP server in the left pane and choose Properties. In the DNS tab (see Figure 3.7), select the Enable DNS dynamic updates according to the settings below check box. Then choose whether to allow the DHCP server to make the updates or to override client settings and always perform updates.

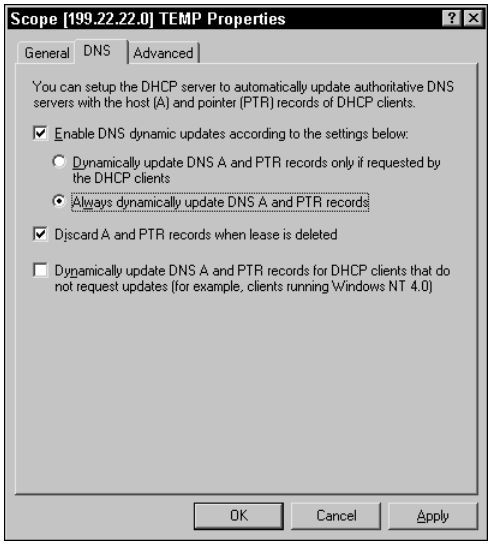

FIGURE 3.7 The Always dynamically update DNS A and PTR records option has been set.

You can also set the DHCP server to update A and PTR records of clients that do not make any dynamic update requests, such as Windows NT 4 systems. On the DNS server, you specify the DHCP server as the only computer authorized to update DNS entries.

If you use multiple Windows Server 2003 DHCP servers on your network and configure your zones to allow secure dynamic updates only, you need to add your DHCP servers to the built-in DnsUpdateProxyGroup to grant all your DHCP servers secure rights to perform updates for your DHCP clients. The default Discretionary Access Control List (DACL) entries on the DNS zones stored in Active Directory are as follows:

- Administrators have Allow settings for the following: Read, Write, Create All Child Objects, Special Permissions
- DnsAdmins have Allow settings for the following: Full Control, Read, Write, Create All Child Objects, Delete Child Objects
- Domain Admins have Allow settings for the following: Full Control, Read, Write, Create All Child Objects, Delete Child Objects
- Enterprise Admins have Allow settings for the following: Full Control, Read, Write, Create All Child Objects, Delete Child Objects
- Enterprise Domain Controllers have Allow settings for the following: Special Permissions
- Pre-Windows 2000 Compatible Access has Allow settings for the following: Special Permissions
- System has Allow settings for the following: Full Control, Read, Write, Create All Child Objects, Delete Child Objects

NOTE

When dealing with clients that dynamically update their own DNS records, host record registration often fails because the primary DNS suffix on the client machine does not match the DNS zone name. For example, say you have gunderville.com for the actual Active Directory domain and DNS domain, and the computer has a primary DNS suffix of 2000trainers.com. This discrepancy causes the system to attempt to register at 2000trainers.com, which it typically is not authorized to do.

Plan a NetBIOS Name Resolution Strategy

WINS is still around in Windows Server 2003 for backward-compatibility with older programs and client systems still used in some enterprises. You can add it to your Windows Server 2003 system through the Configure Your Server Wizard or by clicking Control Panel, Add or Remove Programs, Add/Remove Windows Components.

Only a couple of new features have been added to WINS (sometimes referred to as a *NetBIOS name server*) since the updates released in Windows 2000 Server. In Windows Server 2003, WINS has new filtering and search functions; to help you locate records, only those records that fit search criteria you have set are shown. Some of the available search filters include record owner, record type, NetBIOS name, and IP address with or without subnet mask.

When administrators are configuring a domain's replication strategy, they can define a list that controls the source of incoming name records during pull replication between WINS servers. They can also accept only name records owned by specific WINS servers during replication and exclude the name records of all other WINS servers.

The Windows Server 2003 DNS service enables administrators to configure WINS servers to look up hostnames not found in the DNS domain namespace. To do this, NetBIOS namespaces managed by WINS are checked, so systems that cannot query WINS directly can still resolve a system NetBIOS name to find the resource in WINS via DNS.

You can use WINS forward lookup resource records with a name query to WINS servers as a way to get further name resolution when machine names aren't found in DNS. The Windows Server 2003 WINS service supports WINS clients running on the following platforms:

- Windows Server 2003 (all versions including 64-bit)
- Windows XP (all versions including 64-bit)
- Windows Me
- Windows 2000 Server
- Windows 2000 Professional
- Windows NT Server

- Windows NT Workstation
- Windows 98
- Windows 95
- Windows for Workgroups
- Microsoft LAN Manager
- MS-DOS clients
- OS/2 clients
- Linux and UNIX clients (with Samba installed)

WINS-enabled clients communicate with a WINS server to do the following:

- Register client names in the WINS database
- Renew client names with the WINS database
- Release client names from the WINS database
- Resolve names by obtaining mappings from the WINS database for user-names, NetBIOS names, DNS names, and IP addresses

Clients that are not WINS enabled can use some WINS services through the use of WINS proxies. *WINS proxies* are WINS clients configured on a specific subnet to act on behalf of other host computers that cannot use WINS directly to resolve NetBIOS names. NetBIOS names are 15 characters long, with a 16th reserved character that uniquely identifies a service the particular system is running.

Mainly WINS is designed to resolve NetBIOS names to IP addresses, but in the past, this was often done via broadcast as well. Depending on system configuration and type, some systems use broadcasts only to resolve NetBIOS names on the network. You can configure a WINS proxy to listen on behalf of these systems and to query WINS for names not resolved by broadcast and send the resolution back to the querying client. Normally, WINS proxies are used only on networks that include NetBIOS broadcast-only (B-node) clients.

Systems that are configured to use WINS are normally configured as a hybrid (H-node) client, meaning they attempt to resolve NetBIOS names via a WINS server and then try a broadcast (B-node) if WINS is unsuccessful. Most systems can be configured to resolve NetBIOS names in one of four modes:

- *Broadcast (B-node)*—Clients use a broadcast only to resolve names. An enhanced B-node setting has the client use an `LMHOST` file as well. The hex value for this setting is `0x1`.

- *Peer-to-Peer (P-node)*—Clients use WINS only to resolve names. The hex value for this setting is `0x2`.
- *Mixed (M-node)*—Clients first use a broadcast in an attempt to resolve NetBIOS names. If this fails, they attempt the resolution via the WINS server. The hex value for this setting is `0x4`.
- *Hybrid (H-node)*—Clients first use the WINS service in an attempt to resolve NetBIOS names. If this fails, they attempt the resolution via broadcast. The hex value for this setting is `0x8`.

Enhanced B-node systems (specifically, others can be configured to use it as well) use the `LMHOSTS` file, a text file that is manually updated with NetBIOS name–to–IP address mapping information on a network. The file is located in the `%SystemRoot%\System32\Drivers\Etc` folder by default on Windows 2000, XP, and Server 2003 systems.

Additional information on `LMHOSTS` files is available in "Need To Know More?" at the end of this chapter. Using `LMHOSTS` files and manually updating them are usually administratively viable only when fewer than 25 systems total are in use. WINS is the best resolution in an environment with multiple subnets.

The following is a quick overview of the straight WINS registration process:

1. Systems receive their IP addresses from a DHCP server, which also configures them to use a specific WINS server or servers.
2. The client sends a name registration request directly to that WINS server.
3. If the system's NetBIOS name is not already in the WINS server's database, it is accepted as a new registration.
4. The name is entered with a new version ID, given a timestamp, and marked with the WINS server's owner ID. The timestamp is the length of time the system can use the name it has registered on the network.
5. A registration response is sent back to the registering client system from the WINS server with a TTL value equal to the timestamp recorded for the name.

The default renewal interval for WINS names is six days. Clients attempt to renew their registrations when 50% of the TTL value has elapsed. A name must be refreshed before this interval ends, or the WINS server force-releases it.

WINS servers support burst handling for name registration during peak traffic times. Name registration bursts normally occur first thing in the morning if a location has a set start time, and a large number of users come in to work at once and simultaneously start up their PCs.

WINS has a normal name registration operating threshold (called a *queue value*) of 500 by default. If registering systems stay below this value, they are handled via the normal registration process. If this queue value is exceeded, the WINS server responds to new client registration requests with a shorter lease threshold. The TTL is set much lower than the default six days (or whatever the current setting is), which forces clients to reregister with WINS.

The first 100 clients over the set queue value are given a TTL setting of five minutes on their names. The next 100 receive a TTL of 10 minutes, and so on, until a maximum of 50 minutes is reached. Then the 100-count process starts over at five minutes. As the load on the WINS server falls below the queue value, clients reregister and receive the full lease time for their NetBIOS names.

A single WINS server can handle up to 10,000 clients for NetBIOS name resolution requests, based on a decent CPU and adequate memory and disk throughput, but other factors need to be considered, such as fault tolerance, server hardware, other services running on the server, network topology, number of remote users, and so on.

WINS Replication

In most cases, especially in larger environments, multiple WINS servers are running. Some clients might be configured to use one WINS server, and other sets of clients might be configured to use another. This causes a problem if a client using WINS1 needs to resolve the name of a system that registered on WINS2. *Convergence time* is the time it takes to replicate a new entry in a WINS database to all the other WINS servers in the environment.

NOTE

When a client registers its name with a WINS server and an entry is made to that database, it is called an *originating update*. When that change is replicated to other WINS databases, it is referred to as a *replicated update*.

When planning placement and replication for WINS servers, you must decide on an acceptable convergence time for your network. Replication is used to synchronize WINS databases across the different WINS servers.

WINS servers are configured to sync with replication partners. Those partner WINS servers can be configured as pull partners, push partners, or a combination push/pull partner. *Pull partners* pull database entries from their WINS replica partners at a predetermined time. This means that no matter how many (or how few) changes have been made to the WINS database, at the time of pull replication, the server gets the updates. *Push partners* work off a quota number. When a WINS server has accrued a predetermined number of updates to its database, it pushes those changes out to its replication partners. In other words, push partners send updates out as soon as they have enough accumulated database updates, regardless of how long it has been since the last updates were sent.

There are issues with both designs. Pull partners, which might be configured to get databases only once every hour, are not properly updated during high-volume registration hours, such as the beginning of the workday. Push partners, which might be configured to replicate only when they have 300 new records, do nothing if the count stops at 153. The solution is to set up *push/pull partners*. These WINS servers are normally paired up, or if there are enough of them, configured in a hub and spoke setup to push and pull their changes among each other, thus solving both problems mentioned previously.

Exam Prep Questions

1. Which OSI layer is responsible for translating logical network address and names, such as computer names, to their MAC addresses and for addressing and routing data packets over the network?

   - ❍ **A.** Transport layer
   - ❍ **B.** Physical layer
   - ❍ **C.** Network layer
   - ❍ **D.** Data Link layer

2. You are the network administrator for Contoso Ltd. There is a single Active Directory domain called contoso.com. All servers in the domain are running Microsoft Windows Server 2003. Client computers are running Microsoft Windows XP Professional.

   The company has a head office and a branch office. The offices are connected via a T1 link. There is a router in each office.

   You need to implement a DHCP solution to provide dynamic IP addressing to DHCP clients in both offices. The solution must minimize network traffic on the WAN connection. There should not be a single point of failure in the DHCP infrastructure. What should you do?

   - ❍ **A.** Place two DHCP servers in the head office and place them in a cluster configuration. Configure the router in the branch office as a DHCP Relay Agent.
   - ❍ **B.** Place two DHCP servers in each office. Configure each set of DHCP servers in a cluster configuration.
   - ❍ **C.** Place two DHCP servers in the head office and place them in a cluster configuration. Configure a computer in the branch office as a DHCP Relay Agent.
   - ❍ **D.** Place two DHCP servers in the head office. Distribute the scopes for both offices across both servers using the 80/20 rule. Configure a computer in the branch office as a DHCP Relay Agent.

3. How is the division point between the network ID and the host ID of an IP address calculated?

   - ❍ **A.** The Internet Engineering Task Force (IETF) sets the standard for calculating the division point.
   - ❍ **B.** The Least Significant Bit (LSB) is used.
   - ❍ **C.** The Most Significant Bit (MSB) is used.
   - ❍ **D.** A subnet mask is used.

4. You are the network administrator for your company. Servers are running Microsoft Windows Server 2003. Client computers have been upgraded to Windows XP Professional.

You are planning the subnet configuration for the network. The network currently needs to support 12 subnets. This number is expected to reach 16 within the next year. You need to choose a subnet mask that supports this configuration. What should you do?

- ❍ **A.** Use a subnet mask of 255.255.255.192.
- ❍ **B.** Use a subnet mask of 255.255.255.224.
- ❍ **C.** Use a subnet mask of 255.255.255.240.
- ❍ **D.** Use a subnet mask of 255.255.255.248

5. Which routing protocols support Classless Inter-Domain Routing (CIDR)? (Choose two.)

- ☐ **A.** RIP version 2
- ☐ **B.** Open Shortest Path First (OSPF)
- ☐ **C.** RIP version 1
- ☐ **D.** Exterior Gateway Protocol
- ☐ **E.** Interior Gateway Protocol

6. You are the network administrator for Contoso Ltd. The network consists of multiple subnets connected by routers.

You have finished installing a Windows Server 2003 DHCP server. You create the necessary scopes and configure the 003 router option to assign all clients the IP address of their local router. All clients successfully lease an IP address. However, you discover that users on Subnet A are the only ones capable of communicating outside their local subnet. What could be causing the problem?

- ❍ **A.** All the scopes have not yet been activated.
- ❍ **B.** The DHCP option is configured at the server level.
- ❍ **C.** The DHCP server has not yet been authorized.
- ❍ **D.** The 003 router option must first be activated.

7. Given the Class C IPv4 address range of 192.199.199.0, can you use a subnet mask of 255.255.255.240 and create at least 17 subnetworks with at least 13 hosts per subnet?

- ❍ **A.** Yes, you can use a subnet mask of 255.255.255.240 and create at least 17 subnetworks with at least 13 hosts per subnet.
- ❍ **B.** No, the subnet mask of 255.255.255.240 is invalid for the Class C version 4 IP address ranges.

- ❍ **C.** 192.199.199.0 is not a Class C address range.
- ❍ **D.** No, there would be fewer than 13 hosts per subnet.
- ❍ **E.** No, there would be fewer than 17 subnetworks.

8. You are the network administrator for your company. All servers are running Microsoft Windows Server 2003. Client computers are running Microsoft Windows 95 and Microsoft Windows 2000 Professional.

To improve accuracy on several WINS servers while conserving available bandwidth, you have decided to replicate WINS servers only at night. Which type of WINS replication can only be configured for an interval?

- ❍ **A.** Push/Pull
- ❍ **B.** Push
- ❍ **C.** Pull
- ❍ **D.** Pull/Push

9. Maximum Transmission Units (MTUs) are based on the type of network that is installed. What is the MTU of ethernet networks?

- ❍ **A.** 1,460 bytes
- ❍ **B.** 1,500 bytes
- ❍ **C.** 4,464 bytes
- ❍ **D.** 17,914 bytes

10. You are the network administrator for Contoso Ltd. The company has upgraded existing servers to Microsoft Windows Server 2003.

Contoso.com is now considering a name for their DNS structure on the inside of their firewall. The main goal is maximum security, even if new training for users is required. Which of these names should you use?

- ❍ **A.** contoso.com
- ❍ **B.** contoso.com.ad
- ❍ **C.** ad.contoso.com
- ❍ **D.** newname.ad

11. How do you configure clients to use APIPA addressing?

- ❍ **A.** APIPA addressing needs to be set manually in the Registry.
- ❍ **B.** The client needs to be configured to obtain an IP address automatically.

- ❍ **C.** The client needs to have the APIPA option selected in the IP Properties dialog box.
- ❍ **D.** The DHCP server needs to have an APIPA scope configured and authorized.

12. You are the network administrator of a Windows Server 2003 network. You have decided to implement a DNS solution for your network, including a small remote office.

The remote office does not have its own domain controller. Your main goal is to minimize all replication and zone transfer traffic on the slow link between your home office and the remote office. Which type of DNS server should you use in the remote office?

- ❍ **A.** Active Directory Integrated
- ❍ **B.** Secondary
- ❍ **C.** Primary
- ❍ **D.** Caching-Only

13. You are the network administrator for your company. All servers are running Microsoft Windows Server 2003. The network consists of multiple subnets. Routers cannot forward the DHCPDISCOVER broadcast messages. You install the DHCP service on a single server. You are testing the configuration from a Windows XP computer on a different subnet.

The computer fails to obtain an IP address from the DHCP server. The client computer configures itself with an APIPA address. What can be done to correct this problem? (Choose three.)

- ☐ **A.** Install an RFC 1452–compliant router
- ☐ **B.** Install an RFC 1542–compliant router
- ☐ **C.** Install DHCP Relay Agent on the client subnet
- ☐ **D.** Install a local DHCP server
- ☐ **E.** Install a DHCP Relay Agent on the DHCP server
- ☐ **F.** Install a DHCP Relay Agent on the DHCP server subnet

14. You are the network administrator for the Contoso Ltd. All servers are running Microsoft Windows Server 2003. Contoso has seven offices located in different parts of the United States. The central office hosts the primary DNS server. All branch office locations have their own DNS servers configured as secondary servers.

The offices are currently connected by slow WAN links, with no plans to upgrade them. The annual budget allows for the addition of a second DNS server at each of the locations. However, you do not want any more traffic generated from zone transfers on the WAN or the local networks. What should you do?

- ❍ **A.** Configure the new servers as Standard primary DNS servers
- ❍ **B.** Configure the new servers as Standard secondary DNS servers
- ❍ **C.** Configure the new servers as Master name servers
- ❍ **D.** Configure the new servers as Caching-only servers

15. What type of DNS resolution involves a query made from a client to a DNS server, and the server returns the best answer it can provide based on its local cache or stored zone data?

- ❍ **A.** Recursive query
- ❍ **B.** Iterative query
- ❍ **C.** Forward query
- ❍ **D.** Cache query

16. You are a network administrator for your company. Servers are running Microsoft Windows Server 2003 and UNIX. Client computers are running Microsoft Windows XP Professional. The Domain Name System (DNS) server for the network is running UNIX and the most recent version of BIND. The BIND server called DNS01 is a primary server for the internal DNS domain named contoso.com.

You run the Active Directory Installation Wizard (dcpromo.exe) on a server named SRV01 to create an Active Directory domain named ad.contoso.com for the network. You allow the wizard to install and configure DNS on the server with default settings.

On DNS01, you delegate the subdomain ad.contoso.com to SRV01. You plan to configure DNS01 as a secondary server for ad.contoso.com. You need to ensure that SRV01 can communicate with DNS01. Your solution must ensure that the servers communicate as efficiently as possible. What should you do?

- ❍ **A.** Disable the BIND secondaries option on SRV01
- ❍ **B.** Configure ad.contoso.com to use dynamic updates instead of secure dynamic updates
- ❍ **C.** Convert the new DNS zone from Active-Directory Integrated to a primary zone
- ❍ **D.** Create a stub zone for contoso.com on SRV01

17. What are some of the benefits of caching-only DNS servers? (Choose three.)

- ☐ **A.** Local sites that use them do not use WAN bandwidth for DNS resolutions.
- ☐ **B.** They hold a read/write copy of the DNS database on the local site.
- ☐ **C.** Replication is configurable for off hours, thus limiting the impact of zone transfer traffic.

- ☐ **D.** They do not produce any zone transfer traffic.
- ☐ **E.** They reduce traffic across a WAN because they attempt to locate information in their cache to resolve local client requests.

18. When clients dynamically update their own DNS records, what type of host record registration failures can occur?

- ❍ **A.** Incorrect parent DNS suffix listed
- ❍ **B.** Incorrect DNS suffix list entry appended
- ❍ **C.** Incorrect secondary DNS suffix listed
- ❍ **D.** Incorrect primary DNS suffix listed

Answers to Exam Prep Questions

1. **The correct answer is C.** The Network layer is responsible for translating logical network address and names, such as computer names, to their MAC addresses and for addressing and routing data packets over the network. The Transport layer adds another connection below the Session layer and helps manage data flow control between nodes on the network. Therefore, answer A is incorrect. The Physical layer defines the interface between the medium and the device. Therefore, answer B is incorrect. The Data Link layer mainly handles error correction, flow control, and communication with the network adapter card. Therefore, answer D is incorrect.
2. **The correct answer is B.** You should place two DHCP servers in each office and place them in cluster configurations. This solution minimizes network traffic on the WAN link and eliminates DHCP as a single point of failure. Answers A, C, and D are incorrect. By placing the only DHCP servers in the head office, DHCP clients in the branch office must rely on the availability of the WAN link to obtain an IP address.
3. **The correct answer is D.** The division point between the network ID and the host ID is called the subnet mask. The subnet mask is used to determine where the network number in an IP address ends and the node number in an IP address begins. Therefore, answers A, B, and C are incorrect.
4. **The correct answer is C.** The subnet mask of 255.255.255.240 allows for a maximum amount of 16 subnets. This meets the requirements for 12 subnets. Answers A and B are incorrect because these subnet masks do not allow for enough subnets. Answer D is incorrect because this subnet mask allows for more subnets than will ever be required.
5. **The correct answers are A and B.** CIDR is supported by RIP version 2 and OSPF routing. Because CIDR supports multiple subnet masks per subnet, it requires routers that support more advanced interior routing protocols, such as RIP version 2 and OSPF. CIDR isn't supported by RIP version 1 or EGP; therefore, answers C, D, and E are incorrect.

6. **The correct answer is B.** Each subnet has its own gateway, so the 003 router option should be configured at the scope level instead of the server level. Answers A and C are incorrect because all clients are successfully leasing IP addresses. Answer D is incorrect because DHCP options do not have to be activated.

7. **The correct answer is E.** Given the Class C IPv4 address 192.199.199.0, you can use a subnet mask of 255.255.255.240 and create 16 subnetworks with 14 hosts per subnet. Because the question called for at least 17 subnetworks with at least 13 hosts per subnet, the only correct answer is E.

8. **The correct answer is C.** WINS replication can be "pulled" on an interval or "pushed" after a set number of changes to the database; therefore, answers A and B are incorrect because they can be set to push after a number of changes and answer D is incorrect because it is not a valid type of WINS replication.

9. **The correct answer is B.** MTUs are based on the type of network that is installed. Ethernet deployments are limited to a 1,500 byte MTU. Therefore, answers A, C, and D are incorrect.

10. **The correct answer is D.** To maximize security, they should use a name on the inside of the firewall that is totally different than their public name; therefore, answers A, B, and C are incorrect because they still contain the "contoso" name.

11. **The correct answer is B.** To enable APIPA on clients, all you need to do is configure the client to use DHCP (obtain an IP address automatically). When the client starts up and cannot contact a DHCP server, it assigns itself an IP address from the reserved 169.254.0.0 range with a subnet mask of 255.255.0.0. No default gateway is used, and systems that use APIPA are not routable. Therefore, answers A, C, and D are incorrect.

12. **The correct answer is D.** Caching-Only DNS servers do not host a zone and therefore do not have any zone transfer traffic. Active Directory Integrated zones can be hosted only on domain controllers. Even if you installed a domain controller at the remote office, you would then create replication traffic for the domain controller; therefore, answer A is incorrect. Primary and Secondary servers host a zone and create transfer traffic; therefore, answers B and C are incorrect.

13. **The correct answers are B, C, and D.** When a DHCP client request for an IP address hits a non-RFC-1542–compliant router (meaning the DHCPDISCOVER broadcast message is not forwarded off the subnet), it fails to receive a response because the DHCP server never receives the DHCPDISCOVER broadcast message and the client system configures itself with an APIPA address. If a DHCP Relay Agent is in use on the subnet, it receives the DHCPDISCOVER broadcast message and routes the message off the subnet to the DHCP server. Subsequently, when the DHCP server responds with an address and the DHCP client selects the IP address, the client responds with a DHCPREQUEST broadcast message, which includes the IP address of the server that had its offer accepted. Again, this DHCPREQUEST broadcast message does not get out of this subnet unless a DHCP Relay Agent is in use on the subnet, and can receive the DHCPREQUEST broadcast message and forward (route) the message off the subnet to the DHCP server.

14. **The correct answer is D.** By configuring caching-only servers within each location, you can decrease the name resolution response time for users. Because the caching-only servers do not maintain any zone information, no traffic is generated from zone transfers. Therefore, answers A, B, and C are incorrect.

15. **The correct answer is B.** Two types of queries can be performed in DNS: iterative and recursive. The situation described in this scenario is an iterative query. A recursive query happens when a client makes a DNS resolution query to a DNS server, and the server assumes the full workload and responsibility for providing a complete answer to the query, therefore, answer A is incorrect. Forward and cache queries are not applicable; therefore, answers C and D are incorrect.

16. **The correct answer is A.** Any BIND server running version 4.9.4 or later supports fast zone transfers. By disabling the BIND secondaries option, SRV01 will perform fast zone transfers with DNS01. Therefore, answers B, C, and D are incorrect.

17. **The correct answers are A, D, and E.** Caching-only DNS servers perform name resolution on behalf of clients and then cache the resulting name resolutions. They are not configured to be authoritative for a DNS zone, and they do not store Standard Primary or Standard Secondary zones locally. Their local cache holds the most frequently requested names and associated IP addresses and are available for use by subsequent client queries.

18. **The correct answer is D.** When dealing with clients that dynamically update their own DNS records, host record registration often fails because the primary DNS suffix listed on the client machine does not match the DNS zone name. For example, the actual Active Directory domain and DNS domain is gunderville.com, but the computer has a primary DNS suffix listed as 2000trainers.com. This causes the system to attempt to register at 2000trainers.com, which it usually is not authorized to do.

Need To Know More?

TCP/IP Frequently Asked Questions: http://www.itprc.com/tcpipfaq/default.htm.

TCP/IP Protocol Suite—Questions and Answers:
http://www.geocities.com/SiliconValley/Vista/8672/network/.

Microsoft Windows—IPv6: http://www.microsoft.com/IPv6 and
http://www.microsoft.com/windowsserver2003/technologies/ipv6/default.mspx.

RFC 2460: Internet Protocol, Version 6 (IPv6) Specification:
http://www.faqs.org/rfcs/rfc2460.html.

TCP/IPv4 Configurable Registry Settings:
http://msdn.microsoft.com/library/default.asp?url=/library/
en-us/wcetcpip/htm/cmconParametersConfigurableUsingRegistryEditor.asp.

Automatic Windows 98/Me TCP/IP Addressing Without a DHCP Server:
http://support.microsoft.com/default.aspx?scid=KB;en-us;q220874.

Planning DHCP Networks:
http://www.microsoft.com/windows2000/en/server/help/default.asp?url=/windows2000/en/server/help/sag_DHCP_imp_PlanningNetworks.htm.

Domain Name System (DNS) Center:
http://www.microsoft.com/windows2000/technologies/communications/dns/.

Domain Name System Security Extensions: http://www.ietf.org/rfc/rfc2535.txt.

Replacing Root Hints with the CACHE.DNS File:
http://support.microsoft.com/default.aspx?scid=KB;EN-US;Q249868&.

How to Write an LMHOSTS File for Domain Validation and Other Name Resolution Issues:
http://support.microsoft.com/?kbid=180094.

Domain Browsing with TCP/IP and LMHOSTS Files: http://support.microsoft.com/?kbid=150800.

LMHOSTS File Information and Predefined Keywords:
http://support.microsoft.com/?kbid=102725.

Verify WINS Registration of Client NetBIOS Names:
http://www.microsoft.com/windows2000/en/server/help/default.asp?url=/windows2000/en/server/help/sag_WINS_pro_VerifyRegistration.htm.

Registering Names in WINS:
http://www.microsoft.com/windows2000/en/server/help/sag_WINS_und_RegisteringNames.htm.

Windows 2000 Server Windows Internet Naming Service (WINS) Overview:
http://www.microsoft.com/windows2000/techinfo/howitworks/communications/nameadrmgmt/wins.asp.