L3HARRIS
FAST. FORWARD.

INTRODUCTION TO NETWORKING

Module 2: Routing and Addressing

Tanya Wilcox Sr. Specialist Technical Trainer

Course Objectives

- Understand IPv4 and IPv6 addressing schemes, including address classes, subnetting, and classless addressing
- Use the ipconfig/ ifconfig/ ip tools to verify the IP Configuration
- Understand the features and functions of IP routers, such as path selection, routing tables, routing algorithms, and routing metrics.
- Identify common ports and protocols.
- Explain the functions of network services.
- Explain the concepts and characteristics of network switching.

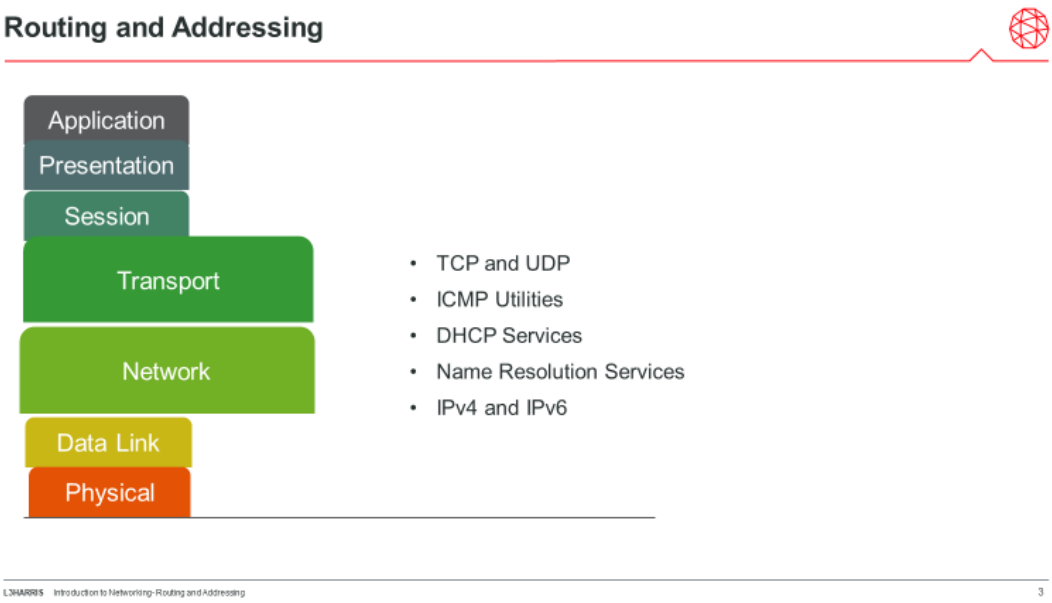

In this lesson, you will identify the addressing and data delivery methods of the Internet Protocol (IP) (Network Layer 3). IP is at the heart of most modern networks, and consequently one of the most important topic areas for a network professional to master.

This lesson will cover the basic format of IPv4 addresses and how they are used to identify networks and hosts. The lesson will also cover IPv6 and methods of assigning an IP address to hosts automatically. We will then continue onto layer 4 and higher to describe the transport and name resolution services that network applications depend upon.

Network Layer (L3) Functions

TCP/IP consists of a suite of complementary protocols and standards that work together to provide the functionality on TCP/IP networks. The Internet Protocol (IP) stands at the heart of this protocol suite, providing logical addressing and packet forwarding between different networks.

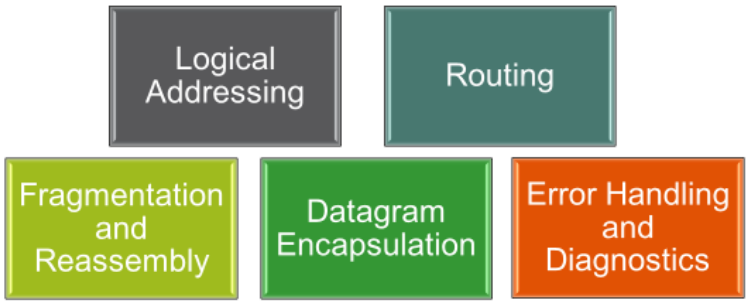

Logical Addressing: Every device that communications over a network has associated with it a logical address, sometimes called layer three address. For example, on the internet, the Internet Protocol (IP) in the network layer protocol and every machine has an IP address Note that addressing is done at the data link layer as well but those addresses refer to local physical devices. In contrast, logical addresses are independent of specific hardware and must be unique across an entire internetwork.

Routing: Moving data across a series of interconnected networks is a defining function of the network layer. It is the job of the devices and software routines that function at the network layer to handle incoming packets from various sources, determine their final destination , and then figure out where they need to be sent to get them where they are supposed to go.

Datagram Encapsulation: The network layer normally encapsulates messages received from higher layers by placing them into datagrams (also called packets) with a network layer header.

Fragmentation and Reassembly: The network layer must send messages down to the data link layer for transmission. Some data link layer technologies have limits on the length of any message that can be sent. If the packet that the network layer wants to send is too large, the network layer must split the packet up, send each piece to the data link layer, and then have pieces reassembled once they arrive at the network layer on the destination machine.

Error handling and diagnostics: Special protocols are used at the network layer to allow devices that are logically connected, or that are trying to route traffic, to exchange information about the status of hosts on the network or the devices themselves.

Network Layer (L3) Common Protocols and Standards

Common Layer 3 Protocols	
IPv4/ IPv6	Internet Protocol
DDP	Datagram Delivery Protocol
ICMP	Internet Control Message Protocol
IGMP	Internet Group Messaging Protocol
GRE	Generic Routing Encapsulation
ESP	Encapsulating Security Payload
AH	Authentication Header
IPSec	Internet Protocol Security
EIGRP	Enhanced Interior Gateway Routing Protocol
OSPF	Open Shortest Path First

Technologies such as Ethernet work at the Physical and Data Link layers of the OSI model (layers 1 and 2). At the Network layer (layer 3), the Internet Protocol (IP) provides logical host and network addressing and routing. IP provides best-effort delivery of an unreliable and connectionless nature. Delivery is not guaranteed, and a packet might be lost, delivered out of sequence, duplicated, or delayed.

There are two versions of IP

- **IPv4** was developed in the 1980's and is still in widespread use today. IPv4 uses a 32- bit address and can support 4.3 billion addresses.
  - Example: *192.168.1.1*
- **IPv6** uses a 128-bit address, introducing a much larger address space of 340 undecillion IP addresses.
  - Example: *2001:0db8:82a3:0000:0000:4a2e:0370:7337*

Protocols/Standards: Some protocols that run directly on IP (rather than over TCP or UDP) include the following:

- Internet Control Message Protocol (ICMP/1) is used for status messaging and connectivity testing.
- Internet Group Messaging Protocol (IGMP/2) is used with multicasting.
- Generic Routing Encapsulation (GRE/47) is used to tunnel packets across an intermediate network. This is used (for example) in some virtual private network (VPN) implementations.
- Encapsulating Security Payload (ESP/50) and Authentication Header (AH/51) are used with the encrypted form of IP (IPSec).
- Enhanced Interior Gateway Routing Protocol (EIGRP/88) and Open Shortest Path First (OSPF/89) are protocols used by routers to exchange information about paths to remote networks

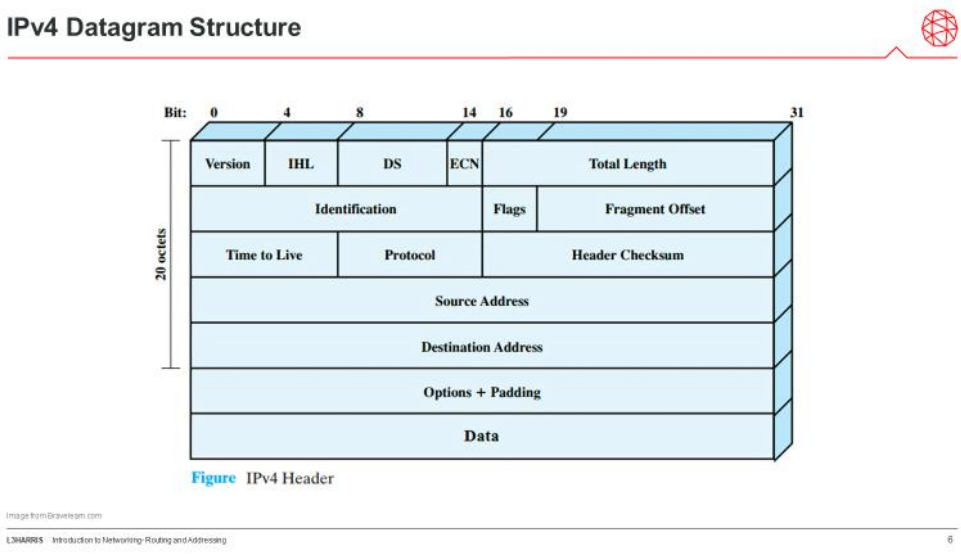

IPv4 HEADER FIELDS The Version field indicates the version of Internet Protocol in use (4), while the Length fields indicate the size of the header and the total packet size (including the payload).

PROTOCOL FIELD The Protocol field describes what is contained (encapsulated) in the payload so that the receiving host knows how to process it.

DIFFSERV FIELD The Differentiated Services Code Point (DSCP) field is used to indicate a priority value for the packet. This can be used with class of service (CoS) and quality of service (QoS) mechanisms to facilitate better quality real-time data transfers, such as video streaming or Voice over IP calling.

TIME TO LIVE FIELD The Time to Live (TTL) is nominally the number of seconds a packet can stay on the network before being discarded; otherwise, packets could endlessly loop around the internet. While TTL is defined as a unit of time (seconds), in practice, it is interpreted as a maximum hop count.

Header Checksum*:* A checksum computed over the header to provide basic protection against corruption in transmission.

ID, FLAGS, AND FRAGMENT OFFSET FIELDS The ID, Flags, and Fragment Offset fields are used to indicate whether the IP datagram has been split between multiple packets for transport over the underlying Data Link protocol.

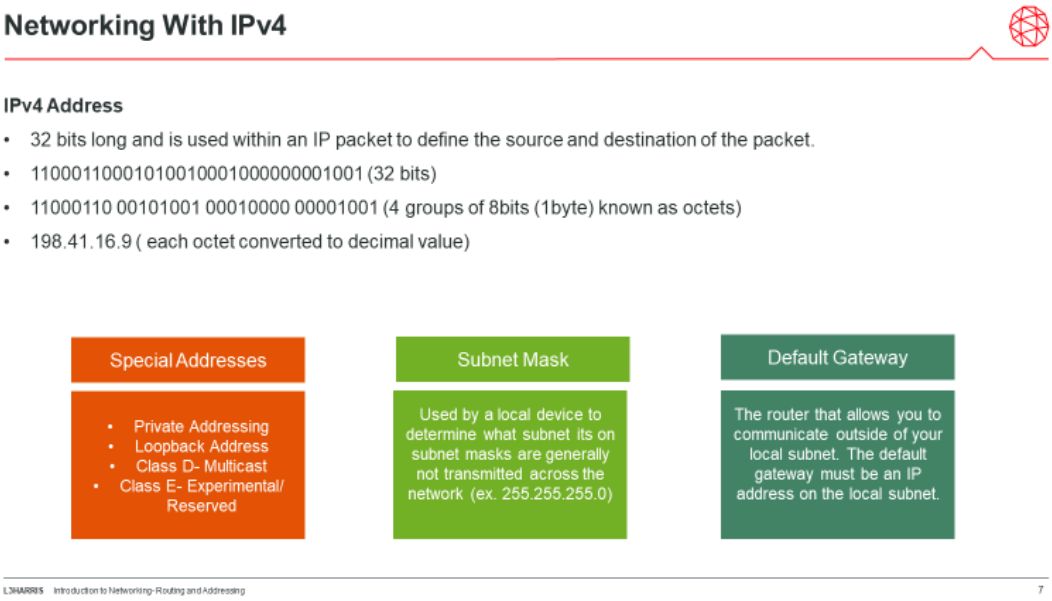

All networks must have a way of uniquely identifying individual computers. This identifier may be in the form of a name or number. At the Data Link layer, each interface is identified by using a MAC or hardware address. This type of address can be used only for local delivery of frames.

At the TCP/IP Internet layer (the OSI Network layer), an IP address is used to identify each host. The IP address provides two pieces of information:

- The network number (network ID)—This number is common to all hosts on the same IP network.
- The host number (host ID)—This number identifies a host within an IP network.

An IPv4 address is 32 bits long and is used within an IP packet to define the source and destination of the packet. In its raw form, it appears as follows:

11000110001010010001000000001001

The 32 bits are subdivided into four groups of 8 bits (1 byte) known as octets. The previous IP address could therefore be written as:

11000110 00101001 00010000 00001001

This representation of an IP address makes human memorizing of the number almost impossible, much less entering it correctly into configuration dialog boxes. To make IP addresses easier to use, they are usually displayed in dotted decimal notation. This notation requires each octet to be converted to a decimal value. The decimal numbers are separated using a period. Converting the previous number to this notation gives: 198.41.16.9

BINARY/DECIMAL CONVERSION

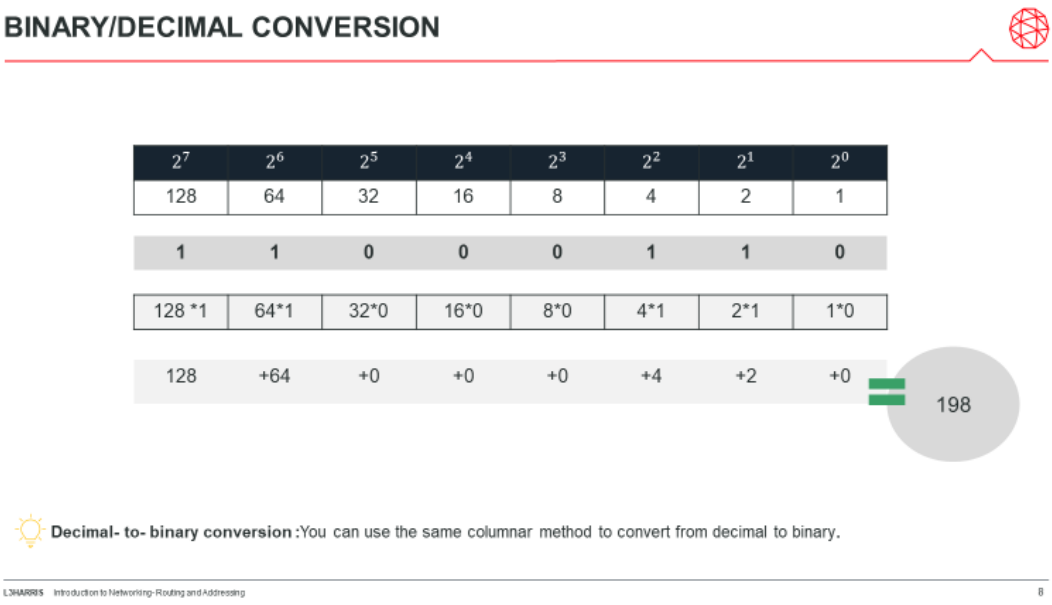

2^7	2^6	2^5	2^4	2^3	2^2	2^1	2^0
128	64	32	16	8	4	2	1
1	**1**	**0**	**0**	**0**	**1**	**1**	**0**
128 *1	64*1	32*0	16*0	8*0	4*1	2*1	1*0
128	+64	+0	+0	+0	+4	+2	+0

198

Decimal- to- binary conversion : You can use the same columnar method to convert from decimal to binary.

Basics of Binary Math:

- A bit= zero or one
  - "off or on" , "cold or hot", "False or true", 0 or 1
- A byte- Eight bits
  - Often referred to as an octet

Binary-to-decimal conversion chart

The base of any number system tells you two things: how many different values any given digit can have and the factor by which the value of a digit increases as you move from right to left in a number.

In base 2 (binary), digits can take one of two different values (0 and 1). The place values are powers of 2 (21=2, 22=4, 23=8, 24=16, and so on). The chart referenced above contains 8 slots, depending on how many bits needed in your conversion, you can continue to extend this out by doubling the last number on the left.

Recall the IP address from the previous slide. 11000110 00101001 00010000 00001001 (198.41.16.9)

Consider the first octet **11000110** represented in base 2. This image shows the octet in the third row, the representation of the place value of each digit of the octet in the fourth row, and the decimal equivalent in the last row. Complete these calculations for each octet to convert the entire IP address.

Note:

If all the bits in an octet are set to 1, the number obtained is 255 (the maximum possible value). Similarly, if all the bits are set to 0, the number obtained is 0 (the minimum possible value). Therefore, theoretically an IPv4 address may be any value between 0.0.0.0 and 255.255.255.255. However, some addresses are not permitted or are reserved for special use

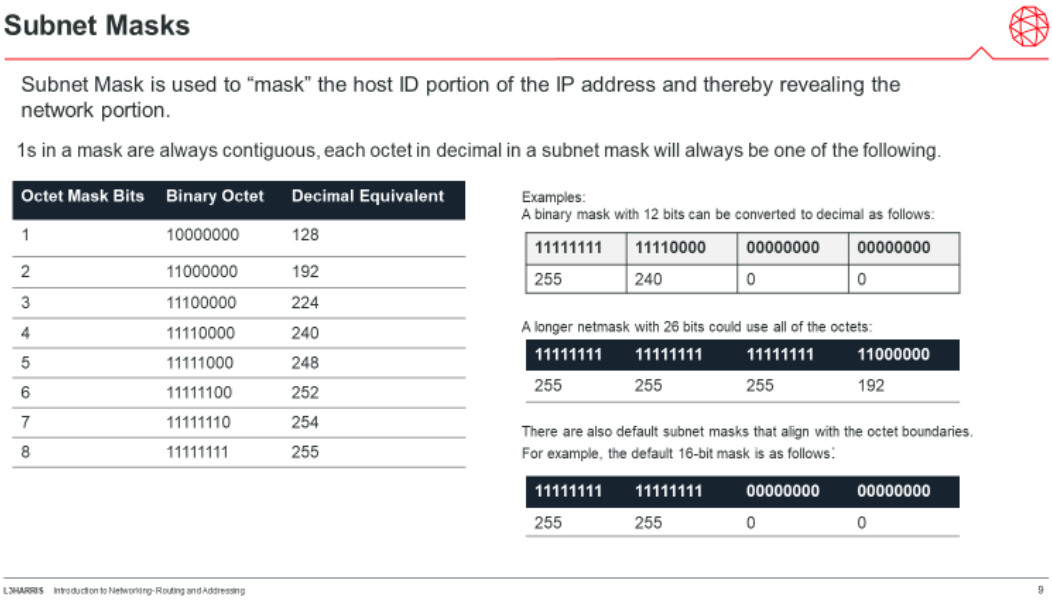

Subnet Masks

Subnet Mask is used to “mask” the host ID portion of the IP address and thereby revealing the network portion.

1s in a mask are always contiguous, each octet in decimal in a subnet mask will always be one of the following.

Octet Mask Bits	Binary Octet	Decimal Equivalent
1	10000000	128
2	11000000	192
3	11100000	224
4	11110000	240
5	11111000	248
6	11111100	252
7	11111110	254
8	11111111	255

Examples:
A binary mask with 12 bits can be converted to decimal as follows:

11111111	11110000	00000000	00000000
255	240	0	0

A longer netmask with 26 bits could use all of the octets:

11111111	11111111	11111111	11000000
255	255	255	192

There are also default subnet masks that align with the octet boundaries.
For example, the default 16-bit mask is as follows:

11111111	11111111	00000000	00000000
255	255	0	0

L3HARRIS Introduction to Networking- Routing and Addressing 9

Subnet Masks: An IP address represents both a network ID and a host ID. A subnet mask (or netmask) is used to distinguish these two components within a single IP address. It is used to "mask" the host ID portion of the IP address and thereby reveal the network ID portion.

- Wherever there is a binary 1 in the mask, the corresponding binary digit in the IP address is part of the network ID.
- The relative sizes of the network and host portions determine how many networks and hosts per network an addressing scheme can support.
- The 1s in the mask are always contiguous.
  - Valid: : 11111111 11110000 00000000 00000000
  - Non Valid: 11111111 00000000 11110000 00000000

A longer netmask with 26 bits could use all the octets:

11111111 11111111 11111111 11000000

255 255 255 192

There are also default subnet masks that align with the octet boundaries. For example, the default 16-bit mask is as follows:

11111111 11111111 00000000 00000000

255 255 0 0

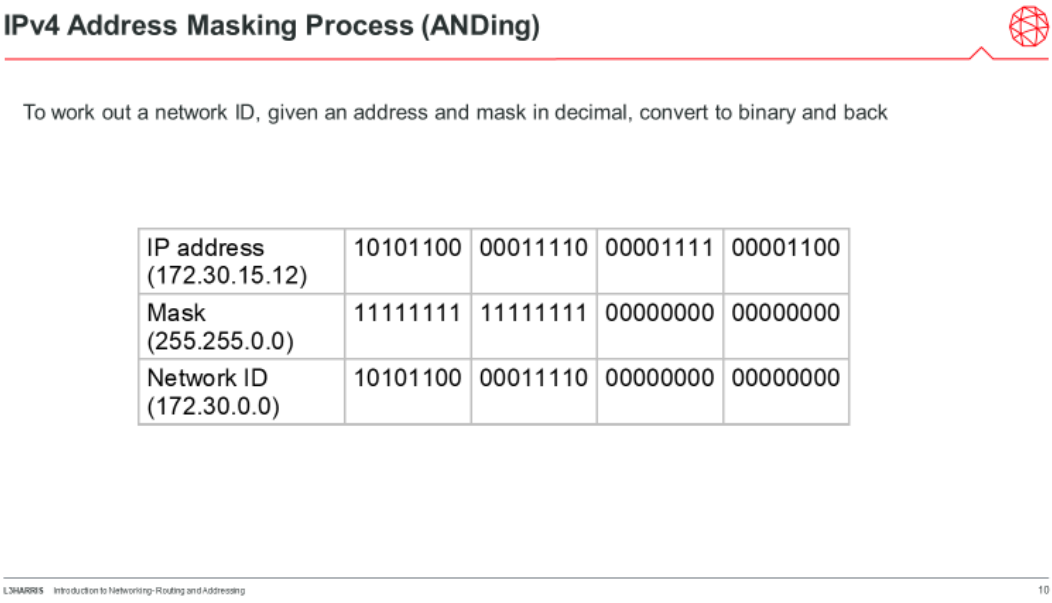

IPv4 ADDRESS MASKING PROCESS (ANDing)

The network ID portion of an IP address is revealed by ANDing the subnet mask to the IP address. The rules for a logical AND are shown in this table.

1 AND 1 = 1

1 AND 0 = 0

0 AND 1 = 0

0 AND 0 = 0

When two 1s are ANDed together, the result is a 1. **Any** other combination produces a 0. For example, to determine the network ID of the IP address 172.30.15.12 with a subnet mask of 255.255.0.0, the dotted decimal notation of the IP address and subnet mask must first be converted to binary notation. The next step is to AND the two binary numbers.

The result can be converted back to dotted decimal notation to show the network ID (172.30.0.0).

172. 30.15.12 => 10101100 00011110 00001111 00001100

255.255. 0. 0 => 11111111 11111111 00000000 00000000

172. 30. 0. 0 => 10101100 00011110 00000000 00000000

The netmask is used by IP to identify whether the source and destination addresses in a packet are on the same IP network. If the destination address has a different network ID, the packet must be sent via one or more routers.

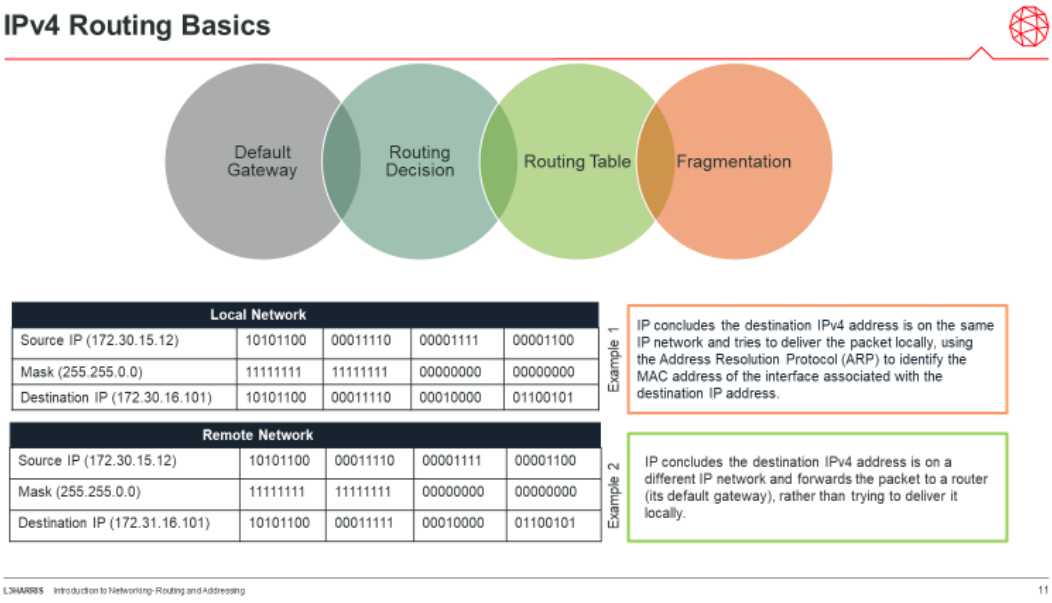

1. IP tries to establish a connection with the destination host by IP address:

- The subnet mask of the host is applied to the source IP address to determine the network address of the source host.
- The subnet mask of the host is applied to the destination IP address to determine the network address of the destination host.

2. The destination network address is compared with that of the source, IP determines if the packet will be routed within or outside of the local network.

3. It is possible that due to limitations in the underlying network, IP may fragment the packet into more manageable pieces (to fit within the MTU of the Data Link protocol frame). If this is the case, IP assigns a new header to each fragment containing:

- A flag to indicate whether more fragments follow.
- A fragment identifier to help group fragments together.
- An offset to assist the destination host in reconstructing the fragments into the original packet.

4. IP then calculates a checksum (to use for error detection) and sends the datagram. A Data Link protocol (such as Ethernet) encapsulates this into one or more frames and transmits them over the network.

5. If the packet has been routed, at the gateway, the Time to Live (TTL) is decreased by at least one. When the TTL is zero, the packet will not be forwarded to another router. This prevents badly addressed packets from permanently circulating the network.

6. The router then determines what to do with the packet by repeating the steps described from the second step on. If the message is destined for yet another network, the process is repeated to take it to the next stage, and so on.

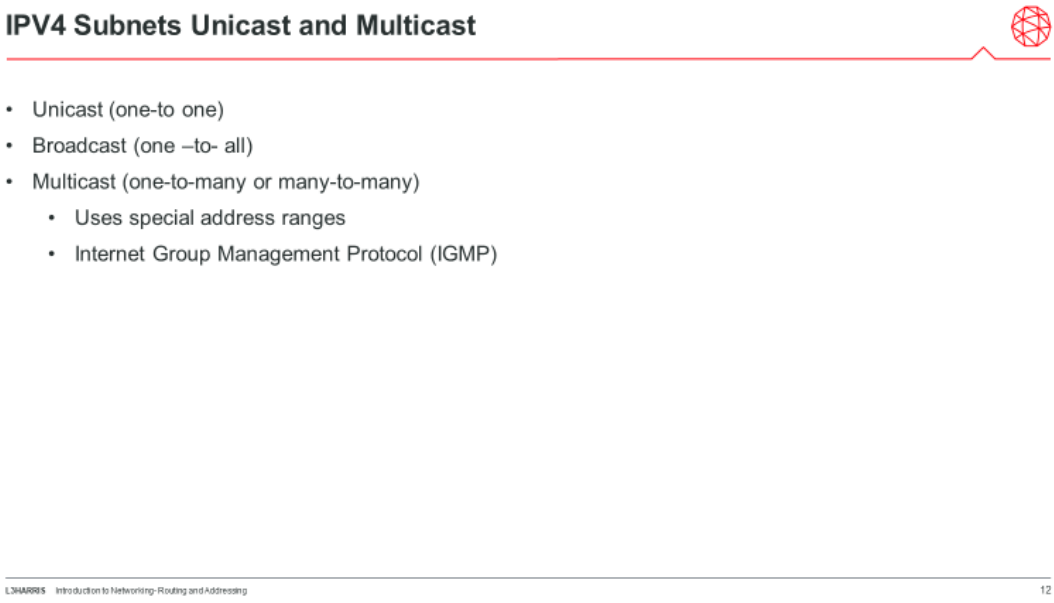

Organizations with large networks need to divide those networks up into smaller segments to improve performance and security. A network segment is represented at the Network layer by a subnet.

- **Unicast**: A packet sent to a single recipient, addressed to the IP address of the destination host.
- **Broadcast**: The destination address is one specially configured to be delivered to all hosts on the local network,
- **Multicast** address :represents a group of computers, programmed to respond to a particular address.

IPv4 multicasting allows one host on the Internet (or private IP network) to send content to other hosts that have identified themselves as interested in receiving the originating host's content (that have joined a multicast group)

.

- Multicast packets are sent to a destination IP address from a special range configured for use with that multicast group.
- The Internet Group Management Protocol (IGMP) is typically used to configure group memberships and IP addresses.
- At layer 2, multicasts are delivered using a special range of MAC addresses. The switch must be multicast capable. If the switch is not multicast-capable, it will treat multicast like a broadcast and flood the multicast transmissions out of all ports.

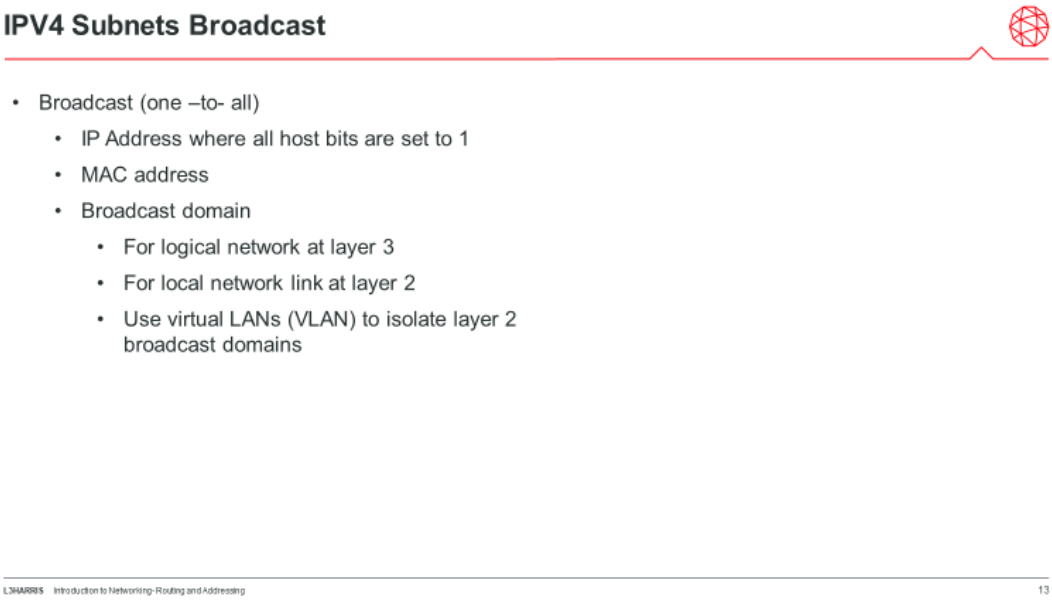

Broadcast Domain: is one where all the hosts receive the same broadcast packets. Boundaries are established at the Network Layer (3) by routers. Routers do not forward broadcasts, except in some specially configured circumstances. Consequently, each IP network is a separate broadcast domain.

The last address in any IP network is the broadcast address, or put another way, the address in any IP network where all the host bits are set to 1.

For example, if the network ID is 192.168.1.0 and the subnet mask is 255.255.255.0, the last octet in the IP address is the host ID portion. If this last octet is set to all 1s, the last address, and therefore the network broadcast address, is 192.168.1.255.

192.168. 1. 0 11000000 10101000 00000001 00000000

255.255.255. 0 11111111 11111111 11111111 00000000

192.168. 1.255 11000000 10101000 00000001 11111111

As with unicast traffic, IP packets must be delivered to hosts using layer 2 MAC addresses. At layer 2, broadcasts are delivered using the group MAC address (ff:ff:ff:ff:ff:ff). This means that there is also a broadcast domain scope at layer 2.

Network efficiency can be achieved by configuring virtual LANs (VLANs) on the switch (or switches).

With **VLANs**, each port is assigned a VLAN ID.

- Each VLAN ID is a separate broadcast domain.
- VLAN IDs can be communicated across multiple switches, which means that users attached to different switches but the same VLAN can be in the same broadcast domain.
- VLANs allow the layer 2 topology to match the layer 3 IP network topology.

IPv4 SUBNET DESIGN

This subnet design allocates:

- Separate subnets (10.0.1.0 and 10.0.2.0) for the two VLANs configured on Switch A and for the serial WAN links configured between Router A and Routers B and C (10.0.3.0 and 10.0.4.0).
- Routers B and C also have a subnet each for their local networks (10.0.5.0 and 10.0.6.0).

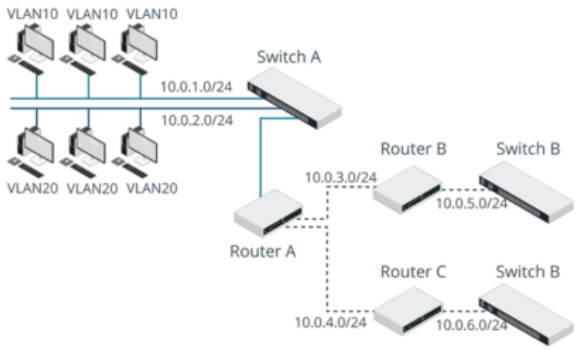

Subnet design. (Image © 123RF.com.)

Subnetting is the process of logically dividing a network into smaller subnetworks (subnets), with each subnet having a unique address.

For example: the referenced subnet design allocates separate subnets (10.0.1.0 and 10.0.2.0) for the two VLANs configured on Switch A and for the serial WAN links configured between Router A and Routers B and C (10.0.3.0 and 10.0.4.0). Routers B and C also have a subnet each for their local networks (10.0.5.0 and 10.0.6.0).

An organization might divide a large IP network into logically distinct subnets for several reasons:

- It is inefficient to have very large numbers of hosts on the same IP network. A single IP network in this sense is a single broadcast domain; excessive broadcast traffic is created when there are many hosts on the same network. Large networks use VLANs to isolate broadcast domains and create subnets to map to each VLAN.
- Networks that use different physical and data link technologies, such as Token Ring and Ethernet, should be logically separated as different subnets.
- Many organizations have more than one site with WAN links between them. The WAN link normally forms a separate subnet.
- It is useful to divide a network into logically distinct zones for security and administrative control.

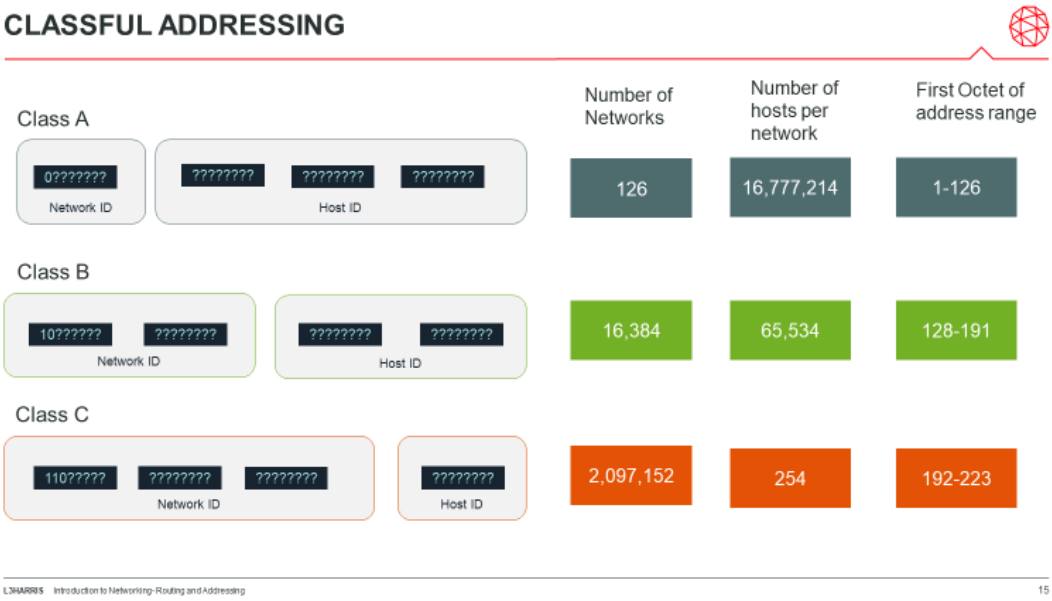

The combination of an IP address and netmask can be used to describe a network ID and a host ID. These parameters allow an internetwork to be divided into logically separate IP networks. Addressing schemes describe different ways of configuring IP addressing to suit different types and sizes of networks.

Classful addressing allocates a network ID based on the first octet of the IP address. The classful addressing scheme was employed in the 1980s, before the use of subnet masks to identify the network ID portion of an address was developed. CLASS A, CLASS B, AND CLASS C ADDRESSES Under classful addressing, the network IDs are divided into three classes, defining different sizes of IP network. When considering classful addressing, you need to identify the address class from the first octet of the IP address

- **Class A** network addresses support large numbers of hosts—over 16 million. However, there are only 126 Class A network addresses.
  - First Octet 1-126
- There are 16 thousand **Class B** networks, each containing up to about 65,000 hosts.
  - First Octet 128-191
- **Class C** networks support only 254 hosts each, but there are over 2 million of them.
  - First Octet 192-223

CLASS D AND CLASS E ADDRESSES

There are two additional classes of IP address (D and E) that use the remaining numbers:

Class D addresses (224.0.0.0 through 239.255.255.255) are used for multicasting.

Class E addresses (240.0.0.0 through 255.255.255.255) are reserved for experimental use and testing.

DEFAULT SUBNET MASKS AND SUBNET IDS

- Aligned along octet boundaries
- Use old class terminology
- Custom subnet masks
  - Modify a default mask
  - Allocate host bits to extra mask bits

Class	Mask
Class A	255.0.0.0
Class B	255.255.0.0
Class C	255.255.255.0

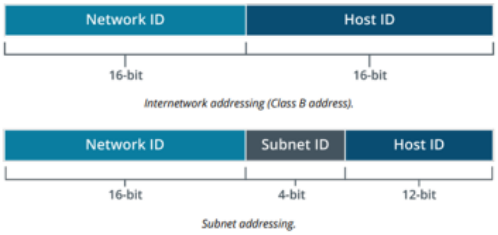

In subnet addressing, the host portion is subdivided into the subnet ID and host ID, so subnet addressing is designed with three hierarchical levels: a network ID, subnet ID, and host ID.

- There is only one subnet mask applied to the IP address on each host. The mask containing the subnet is only used within an IP network.
- External IP networks continue to address the whole network by its network ID.

Using these default masks as examples, you can see how they can be modified to allow a single IP network to be divided into several subnets. To do this, additional bits of the IP address must be allocated as a subnet work address, rather than part of the host ID. The whole network is still referred to by the network ID (by routers external to the network) and the default mask; 172.30.0.0/255.255.0.0 for example. However, routers and hosts within the network add bits to the mask to differentiate the subnets.

For example, if the network designer added 4 bits to the mask, it would mean a subnet mask of 20 bits—the 16 bits of the default Class B mask plus the 4 bits you added. In dotted decimal, the mask would be 255.255.240.0.

172. 30. 0. 0 10101100 00011110 00000000 00000000

255.255. 0. 0 11111111 11111111 00000000 00000000

255.255.240. 0 11111111 11111111 11110000 00000000

This leaves fewer bits (12) available for host IDs, but the purpose of subnetting is to create segments with fewer hosts.

- Note: Wherever a 1 appears in the binary mask, the corresponding digit in the IP address is part of the network or subnet address. Allocate more bits in the mask if you need more subnets. Allocate fewer bits in the mask if you need more hosts per subnet.

Classless Addressing

Classless Inter-Domain Routing (CIDR) is based on **variable-length subnet masking** (**VLSM**), in which network prefixes have variable length (as opposed to the fixed-length prefixing of the previous classful network design).

Benefits of CIDR include:

- Expanded IP address allocations.
- Balanced use of IP address ranges.
- More efficient routing.

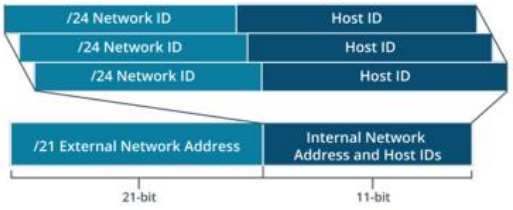

Classless Inter-Domain Routing (CIDR).

With a classless addressing scheme, the concept of address classes and default masks is abandoned in favor of representing the address with an appropriately sized network prefix. CIDR notationin which an IP address is followed by a suffix indicating the number of bits of the prefix.

For example, when expressed in binary, the subnet mask 255.255.240.0 contains 20 ones followed by 12 zeroes. Therefore, the network prefix, expressed in slash notation, is 172.30.0.0/20.

- Note: Most configuration dialog boxes require you to input a subnet mask in dotted decimal format. Some may require you to enter the network address and prefix in **slash notation**, however.

While routers have performed classless routing for years, the class terminology is still very widely used. Even under classless addressing, the old classes are often used as names for the netmasks that align to whole octet boundaries; a Class A network is /8, a Class B network is /16, and a Class C network is /24.

- **More IP address allocations.** Today, we know IPv6 is our long-term IP address solution to the IP address exhaustion problem. However, IPv6 is not yet widely used. In the early 1990s, it was clear we would rapidly exhaust the IPv4 address space if nothing changed. As a result, classless addressing was used as a medium-term solution to help us stretch the life of IPv4.
- **More balanced use of IP address ranges.** Classless addressing decoupled the relationship between network size and IP address and allowed for balanced use across what used to be the Class A, B, and C ranges. Far less wasted addresses.
- **More efficient routing.** VLSM and subnetting make route aggregation and classless routing protocols possible. With route aggregation (sometimes called route summarization or supernetting), routing tables can be smaller, reducing resource consumption on routers, and saving bandwidth. Additionally, including network masks in routing protocols allows for more specific routes to be advertised. For example, 198.51.100.0/29 tells us more than 198.51.100.0 (with an implicit /24).

Private and Public IPv4 Addressing Schemes

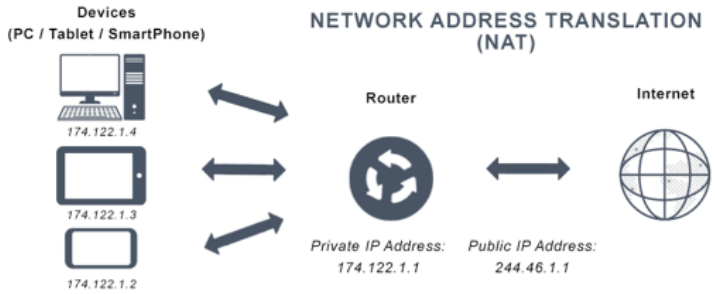

PRIVATE VS. PUBLIC ADDRESSING A public IP network or host address is one that can establish a connection with other public IP networks and hosts over the Internet. The allocation of public IP addresses is governed by IANA and administered by regional registries and Internet Service Providers (ISPs). Hosts communicating with one another over a local area network (LAN) could use a public addressing scheme but will more typically use private addressing.

Private IP addresses can be drawn from one of the pools of addresses defined in RFC 1918 as non-routable over the Internet:

- 10.0.0.0 to 10.255.255.255 (Class A private address range).
- 172.16.0.0 to 172.31.255.255 (Class B private address range).
- 192.168.0.0 to 192.168.255.255 (Class C private address range).

Any organization can use private addresses on its networks without applying to a registry or ISP, and multiple organizations can use these ranges simultaneously. Internet access can be facilitated for hosts using a private addressing scheme in two ways

- Through a router configured with a single or block of valid public IP addresses; the router translates between the private and public addresses using a process called Network Address Translation (NAT).

Through a proxy server that fulfills requests for Internet resources on behalf of clients. The proxy server itself must be configured with a public IP address on the external-facing interface.

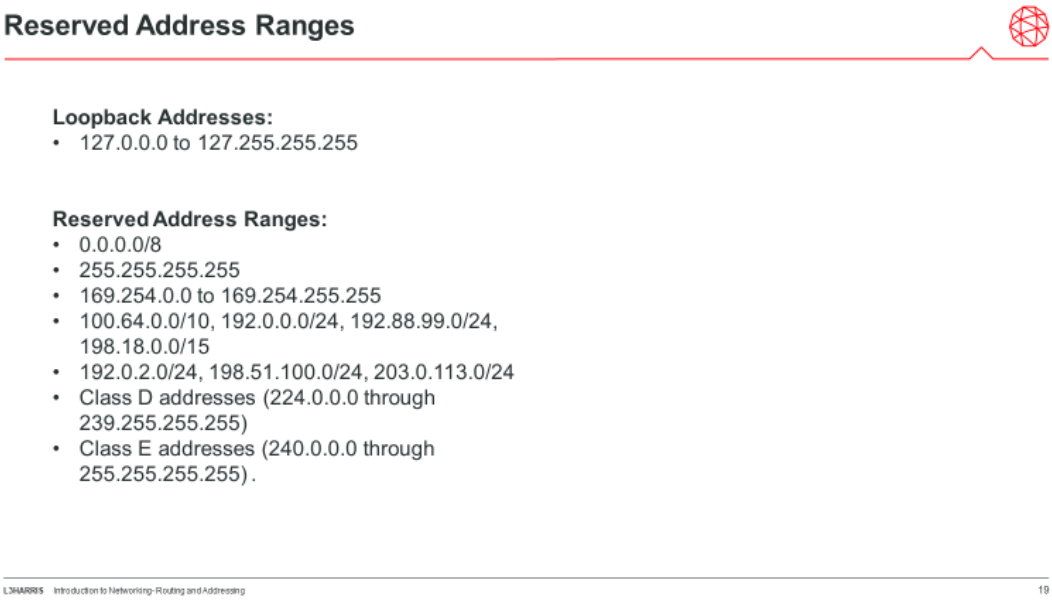

LOOPBACK ADDRESSES: While nominally part of Class A, the range 127.0.0.0 to 127.255.255.255 (or 127.0.0.0/8) is reserved. This range is used to configure a loopback address, which is a special address typically used to check that TCP/IP is correctly installed on the local host. Every IP host is automatically configured with a default loopback address, typically **127.0.0.1**. On some hosts, such as routers, more than one loopback address might be configured. Loopback interfaces can also be configured with an address from any suitable IP range, as long as it is unique on the network (again, often of use in routing).

There are two additional classes of IP address (D and E) that use the remaining numbers:

- Class D addresses (224.0.0.0 through 239.255.255.255) are used for multicasting.
- Class E addresses (240.0.0.0 through 255.255.255.255) are reserved for experimental use and testing.

RESERVED ADDRESS RANGES Additional ranges are reserved for special use and are not publicly routable:

- 0.0.0.0/8—Used when a specific address is unknown. This is typically used as a source address by a client seeking a DHCP lease.
- 255.255.255.255—Used to broadcast to the local network when the local network address is not known.
- 169.254.0.0 to 169.254.255.255—Used by hosts for automatic private IP addressing (APIPA or link-local addressing).
- 100.64.0.0/10, 192.0.0.0/24, 192.88.99.0/24, 198.18.0.0/15—Set aside for a variety of special purposes.
- 192.0.2.0/24, 198.51.100.0/24, 203.0.113.0/24—Set aside for use in documentation and examples.

Summary IPv4

- An IPv4 address is expressed using dotted decimal notation. The netmask can either be expressed using dotted decimal or as a network prefix (slash notation).
- Internet Protocol provides layer 3 addressing, allowing for logically distinct networks
- IPv4 clients are configured with a 32-bit IP address and subnet mask—The subnet mask defines the network and host ID portions of the IP address.
- Hosts with the same network ID transmit packets locally (using ARP). A packet addressed to a host with a different network ID must be sent via a router.
- IPv4 traffic can be addressed to unicast, broadcast, or multicast.
- Classful addressing uses a system of fixed network IDs based on the first octet of the IP address.
- Classless addressing, where the mask is defined by a network prefix, allows for subnetting and supernetting (CIDR), plus variable length subnet masks (VLSMs).
- IPv4 provides public and private addressing schemes. Privately addressed hosts can use some type of NAT or proxy to communicate over the Internet.

IPv6

- 128- bit addressing scheme
- Hexadecimal notation
- Base 16

IPv6 Address Binary Example:

0010 0000 0000 0001 : 0000 1101 1011 1000 :
0000 0000 0000 0000 : 0000 0000 0000 0000 :
0000 1010 1011 1100 : 0000 0000 0000 0000 :
1101 1110 1111 0000 : 0001 0010 0011 0100

- Represented in hex notation as:
  2001:0db8:0000:0000:0abc:0000:def0:1234
- Simplified:
  2001:db8::abc:0:def0:1234

Decimal	Hexadecimal	Binary	Decimal	Hexadecimal	Binary
0	0	0000	8	8	100
1	1	0001	9	9	1001
2	2	0010	10	A	1010
3	3	0011	11	B	1011
4	4	0100	12	C	1100
5	5	0101	13	D	1101
6	6	0110	14	E	1110
7	7	0111	15	F	1111

Conversion Table: Decimal, binary, and hexadecimal values.

IP version 6 (IPv6): Is the most recent version of Internet Protocol developed by Internet Engineering Task Force (IETF) to provide a long-term solution to the problem of address space exhaustion. is base 16 with the possible values of each digit represented by the numerals 0 through 9 and the characters A, B, C, D, E, and F.

- 128-bit addressing scheme has space for 340 undecillion unique addresses.
- Hexadecimal notation is base 16 with the possible values of each digit represented by the numerals 0 through 9 and the characters A, B, C, D, E, and F.

Hexadecimal Numbering: To interpret IPv6 addresses, you must understand hexadecimal notation and the concept of base numbering systems. To start with the familiar; decimal numbering is also referred to as base 10. Base 10 means that each digit can have one of ten possible values (0 through 9). A digit positioned to the left of another has 10 times the value of the digit to the right.

- For example, the number 255 can be written out as follows: (2x10x10)+(5x10)+5 Binary is base 2, so a digit in any given position can only have one of two values (0 or 1), and each place position is the next power of 2.
- The binary value 11111111 can be converted to the decimal value 255 by the following sum: (1x2x2x2x2x2x2x2)+(1x2x2x2x2x2x2)+(1x2x2x2x2x2)+(1x2x2x2x2)+(1x2x2 x2)+(1x2x2)+(1x2)+1

IPv6 Address Structure: IPv6 addresses contain eight 16-bit numbers (double-byte or double-octet), with each double-byte number expressed as 4 hex digits.

- Where a double byte contains leading 0s, they can be ignored. In addition, one contiguous series of 0s can be replaced by a double colon place marker.
- Where IPv6 addresses are used as part of a URL (web address), because both formats use colon delimiters to mean different things, the IPv6 address must be contained within brackets.

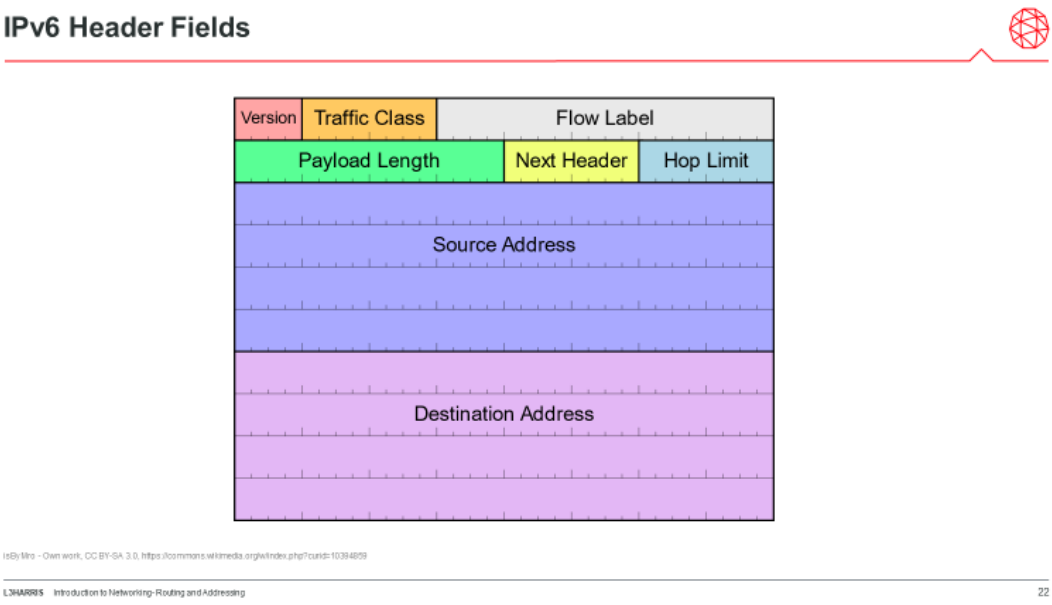

An IPv6 packet consists of two or three elements: the main header, which is a fixed length (unlike in IPv4), one or more optional extension headers, and the payload. The IPv6 packet format is detailed in RFC 2460, but the key features are:

- Version: 4 bits- Used to indicate which version of IP is being used (0110 or 0x06 for IPv6).
- Traffic Class: 8 bits- Describes the packet's priority.
- Flow Label: 20 bits- Used for QoS management, such as for real-time streams. This is set to 0 for packets not part of any delivery sequence or structure.
- Payload Length: 16 bits- Indicates the length of the packet payload, up to a maximum of 64 KB; if the payload is bigger than that, this field is 0 and a special Jumbo Payload (4 GB) option is established.
- Next Header: 8 bits- Used to describe what the next extension header (if any) is, or where the actual payload begins.
- Hop Limit: 8 bits- Replaces the TTL field in IPv4, but performs the same function.
- Source Address: 128 bits- The originating address.
- Destination Address: 128 bits- The target address.

Extension headers replace the Options field in IPv4. There are several pre-defined extension headers to cover functions such as fragmentation and reassembly, security (IPSec), source routing, and so on.

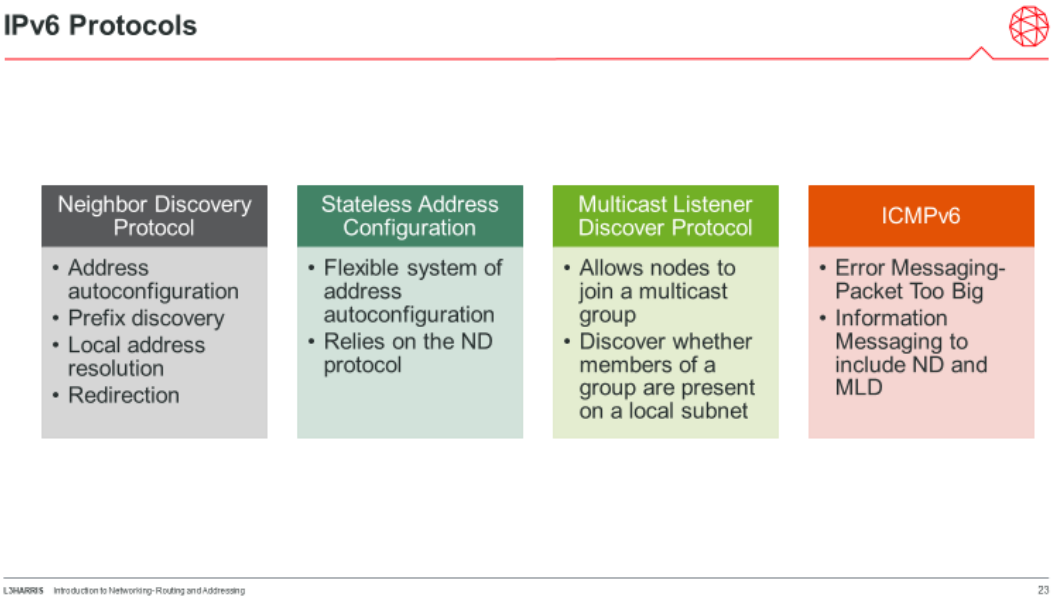

Neighbor Discovery Protocol (ND): performs some of the functions on an IPv6 network that ARP and ICMP perform under IPv4. The main functions are as follows:

- Address autoconfiguration—Enables a host to configure IPv6 addresses for its interfaces automatically and detect whether an address is already in use on the local network, by using neighbor solicitation (NS) and neighbor advertisement (NA) messages.
- Prefix discovery—Enables a host to discover the known network prefixes that have been allocated to the local segment. This also allows next-hop determination (whether a packet should be addressed to a local host or a router). Prefix discovery uses router solicitation **(RS)** and router advertisement **(RA)** messages.
- Local address resolution—Allows a host to discover other nodes and routers on the local network (neighbors). This process also uses neighbor solicitation (NS) and neighbor advertisement (NA) messages.
- Redirection—Enables a router to inform a host of a better route to a particular destination.

Stateless Address Autoconfiguration (SLAAC): Flexible system of address autoconfiguration

- The host generates a link-local address and tests that it is unique by using the Neighbor Discovery (ND) protocol.
- The host listens for a router advertisement (RA) or transmits a router solicitation (RS) using ND protocol messaging. Routers send out advertisements periodically and will respond to a solicitation request to enable clients to determine in which network they reside.

Multicast Listener Discovery Protocol (MLD): Allows nodes to join a multicast group and discover whether members of a group are present on a local subnet.

ICMPv6: updated version of ICMP to include error messaging and informational messaging.

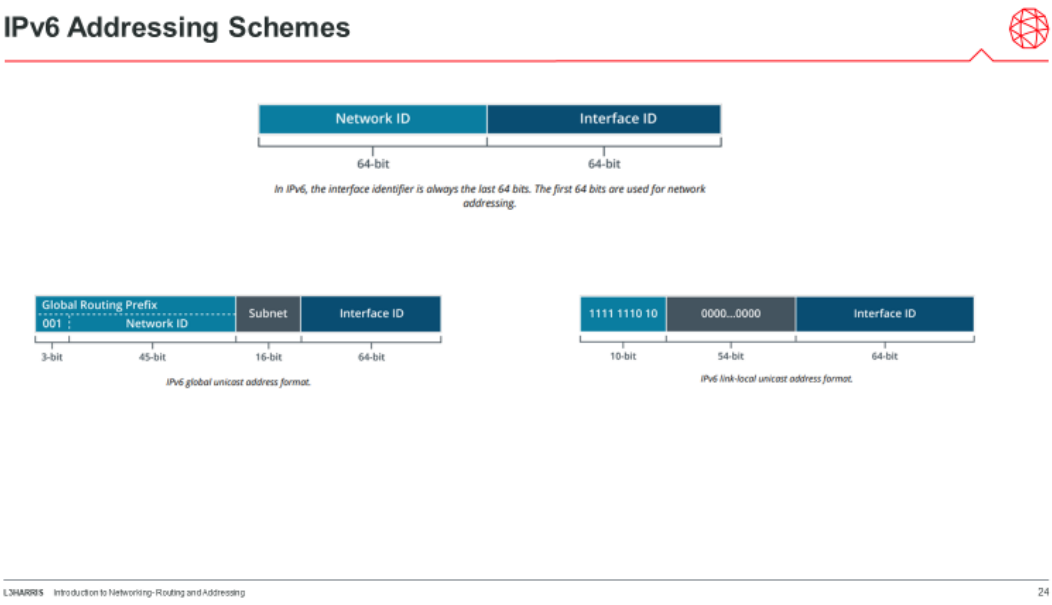

An IPv6 address is divided into two parts: the first 64 bits are used as a network ID, while the second 64 bits designate a specific interface. Network addresses are written using CIDR notation, where /nn is the length of the routing prefix in bits. Within the 64-bit network ID, as with CIDR, the length of any given network prefix is used to determine whether two addresses belong to the same IP network.

For example, if the prefix is /48, then if the first 48 bits of an IPv6 address were the same as another address, the two would belong to the same IP network. This means that a given organization's network can be represented by a network prefix 48 bits long, and they then have 16 bits left in the network ID to subnet their network.,

- 2001:db8:3c4d::/48 would represent a network address,
- while: 2001:db8:3c4d:01::/64 would represent a subnet within that network address.

IPv6 Unicast: IPv6 unicast addressing is scoped; a scope is a region of the network. **Global scopes** provide the equivalent of public addressing schemes in IPv4, while **link-local** schemes provide private addressing.

- **Globally scoped** unicast addresses are routable over the Internet and are the equivalent of public IPv4 addresses.
- **Link-local addresses** are used by IPv6 for network housekeeping traffic. Link-local addresses span a single subnet (they are not forwarded by routers). Nodes on the same link are referred to as neighbors.
- **Unique Local Addressing**: Routable within a site, never routable over the internet.

IPv6 Multicast: Identifies multiple network interfaces. Unlike IPv4, IPv6 routers must support multicast. Broadcast addresses are not implemented in IPv6. Instead, hosts use an appropriate multicast address for a given situation. The well-known multicast addresses are ones reserved for these types of broadcast functionality.

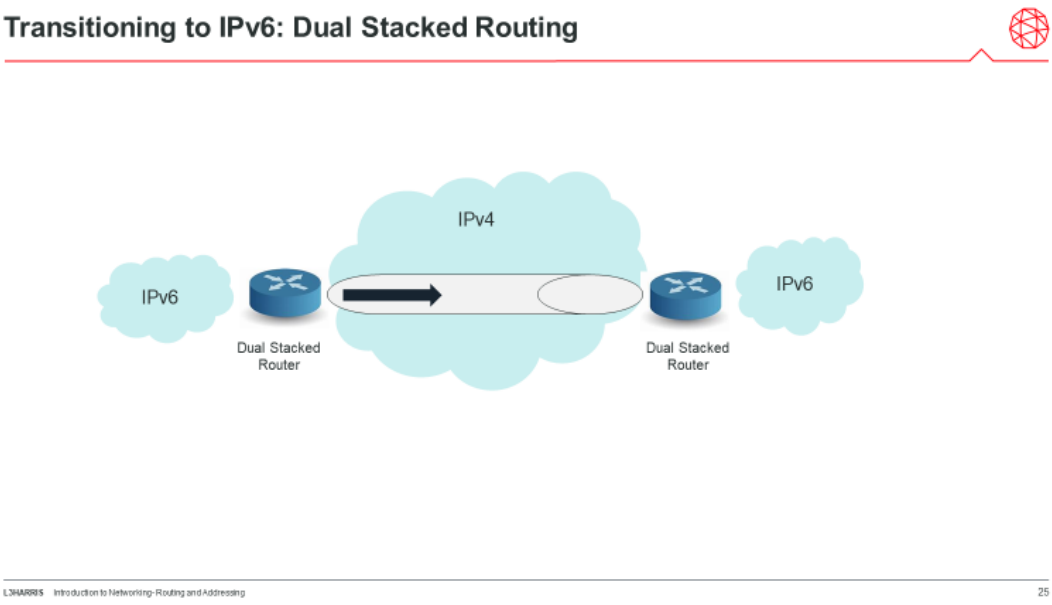

Given the number of devices currently running IPv4, switching to IPv6 is not going to be simple. However, there are two strategies to help make the transition easier.

Dual-stack hosts can run both IPv4 and IPv6 simultaneously and communicate with devices configured with either type of address. Most modern desktop and server operating systems implement dual-stack IP. A dual-stack router can translate between IPv6 and IPv4.

- One technology is Intra-Site Automatic Tunnel Addressing Protocol (ISATAP). Under ISATAP, a dual-stack router takes an IPv6 packet and rewrites it as an IPv4 packet. The ISATAP router allows for a network with mixed IPv4 and IPv6 hosts, but it cannot be used for routing between networks. ISATAP hosts use the link-local range fe80::5efe:w.x.y.z, where w.x.y.z is a dotted decimal IPv4 address.

Dual-stack hosts may also make use of IPv4 mapped addresses.

An IPv4 mapped address is expressed ::ffff:192.168.0.1 (that is, 80 0s followed by 16 1s and then the 32-bit IPv4 address, expressed by convention in dotted decimal). This sort of address is never assigned to hosts, but it is used by IPv4/IPv6 routers to forward traffic between IPv4 and IPv6 networks.

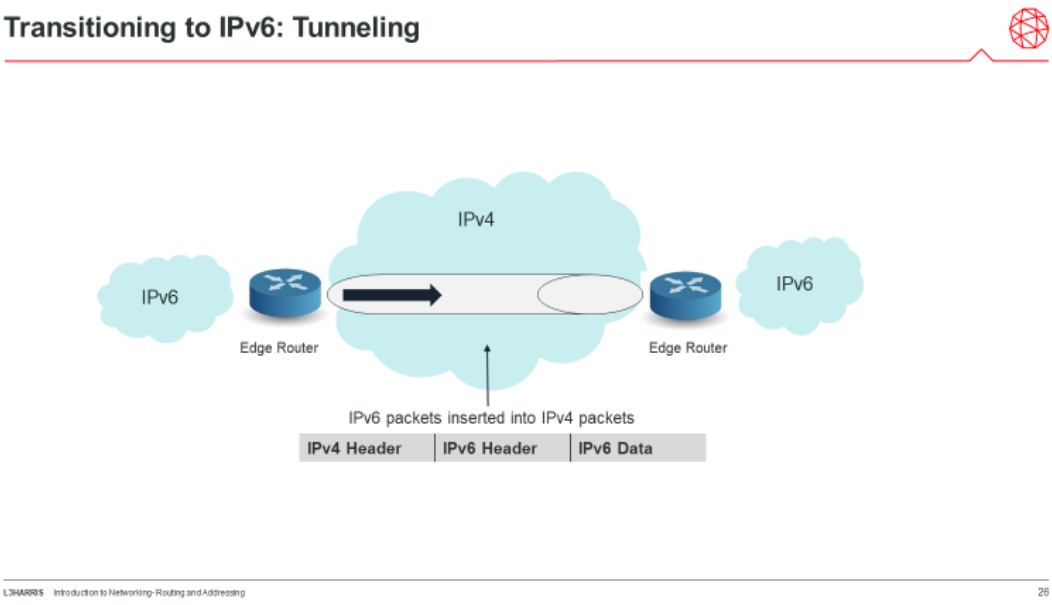

As an alternative to dual-stack routing, tunneling can be used to deliver IPv6 packets across the IPv4 Internet.

- In **6to4** automatic tunneling (RFC 3056), no host configuration is necessary to enable the tunnel. 6to4 addresses start with a leading 2002. Essentially, when 6to4 is implemented, the IPv6 packets are inserted into IPv4 packets and routed over the IPv4 network to their destination. Routing decisions are based on the IPv4 address until the packets approach their destinations, at which point the IPv6 packets are stripped from their IPv4 carrier packets and forwarded according to IPv6 routing rules. 6to4 supports only public IPv4 addresses (that is, those not behind a NAT device).
- Microsoft also provides support for **Teredo** tunneling by Windows hosts. Teredo tunnels IPv6 packets as IPv4-based UDP messages over port 3544. Using UDP rather than TCP allows tunneling through NAT devices. A compatible open-source implementation of Teredo, known as Miredo, is available for UNIX/Linux operating systems.
- **Generic Routing Encapsulation (GRE**). GRE was developed by Cisco and is supported by many Linux distributions and by Microsoft since Windows Server 2012 R2 (with hotfixes). GRE allows a wide variety of Network layer protocols to be encapsulated inside virtual point-to-point links. This protocol has the advantage that because it was originally designed for IPv4, it is considered a mature mechanism and can carry both v4 and v6 packets over an IPv4 network. GRE also has the advantage (like Teredo) of not requiring public IPv6 addresses. It is also possible to tunnel IPv4 through an IPv6 network, in which case the process is known as 4to6 or 4in6 tunneling, as defined in RFC 2473. However, given that the most likely transit network for tunneling between sites is the Internet, which is based on IPv4, this type of tunnel is currently of limited use

Summary IPv6

- IPv6 uses 128-bit addresses with variable length network prefixes and a 64-bit host ID derived from the MAC address or randomly assigned.
- IPv6 supports global and link-local address schemes and unicast and multicast addressing.
- In IPv6, SLAAC, ND, MLD, ICMPv6, and DHCPv6 are used for autoconfiguration and neighbor discovery.
- Communications can take place between IPv4 and IPv6 by using dual-stack hosts or tunneling.

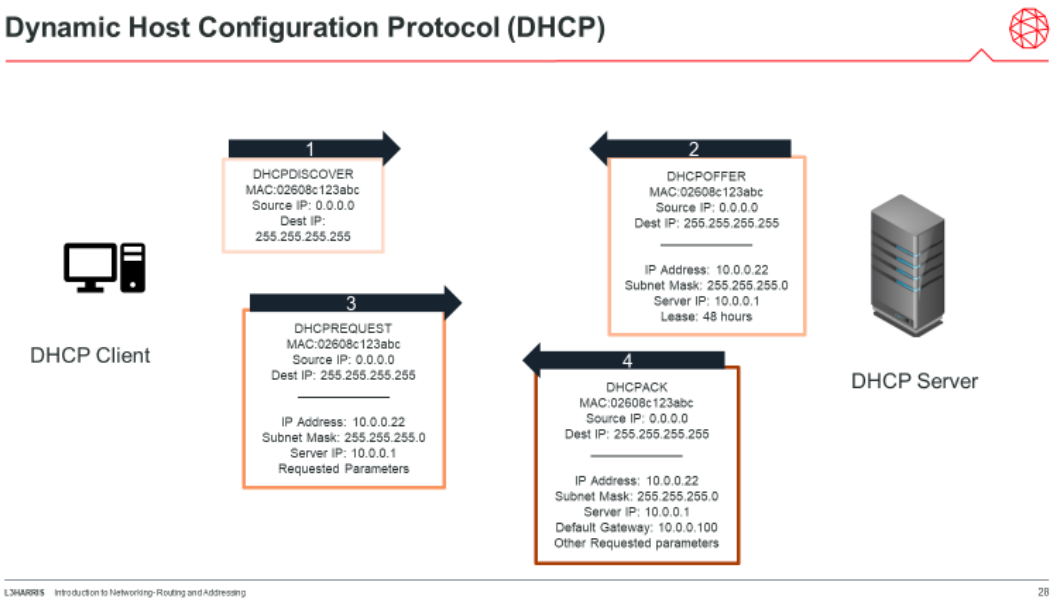

Dynamic Host Configuration Protocol (DHCP): provides an automatic method for allocating an IP address, subnet mask, and optional parameters, such as the Default gateway and DNS server addresses.

- All major Operating Systems provide support for DHCP clients and servers.
- A host is configured to use DHCP by specifying in the TCP/IP configuration that it should automatically obtain an IP address.
  - 1. When a DHCP client initializes, it broadcasts to find a DHCP server. This is called a DHCPDISCOVER packet. All communications are sent using UDP, with the server listening on port 67 and the client on port 68.
  - 2. The DHCP server responds to the client with an IP address and other configuration information, as long as it has an appropriate IP address available. The IP addressing information is offered for a period of time. This packet is also broadcast and is called a DHCPOFFER.
  - 3. The client may choose to accept the offer using a DHCPREQUEST packet—also broadcast onto the network.
  - 4. Assuming the offer is still available, the server will respond with a DHCPACK packet. The client broadcasts an ARP message to check that the address is unused. If so, it will start to use the address and options; if not, it declines the address and requests a new one.

The address is leased, and, after a designated period, the client must theoretically release the IP addressing information. This process does not normally take place since the client can renew or rebind the lease

Automatic Private IP Addressing (APIPA): was developed by Microsoft as a means for clients that could not contact a DHCP server to communicate on the local network anyway.

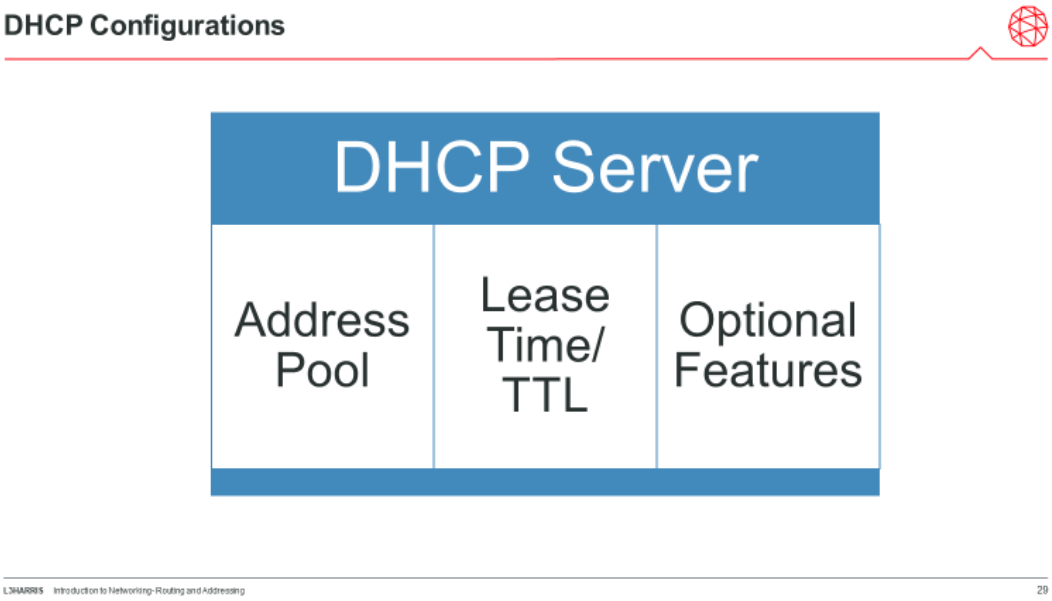

DHCP Server Configuration: DHCP is normally provided as part of the network operating system or through an appliance such as a switch or router. A DHCP server must be allocated a static IP address and configured with the following information:

DHCP Address Pool: The address pool is the range of IP addresses that a DHCP server can allocate to clients on a particular subnet. To define an address pool, you must provide a start and end IP address along with a subnet mask. The subnet mask given must be such that the entire range of addresses is contained within the scope of a single subnet. The server maintains a one-to-one mapping of pools to subnets. That is, no pool can cover more than one subnet and no subnet can contain more than one pool.

DHCP Lease Time/TTL: A lease time defined in the DHCP server for client use of the assigned IP address. The client can renew the lease when at least half the lease's Time to Live (TTL) period has elapsed (T1 timer) so that it keeps the same IP addressing information. If the original DHCP server does not respond to the request to renew the lease, the client attempts to rebind the same lease configuration with any available DHCP server. By default, this happens after 87.5% of the lease duration is up (T2 timer). If this fails, the client reinitializes and continues to broadcast to discover a server.

DHCP Options: Optional features are available depending on the configuration of the network. Each option is identified by a tag byte or decimal value between 0 and 255. Some widely used options include:

- The default gateway (IP address of the router).
- The IP address(es) of DNS servers.
- The DNS suffix (domain name) to be used by the client.
- Other useful server options, such as time synchronization (NTP), file transfer (TFTP), or VoIP proxy.

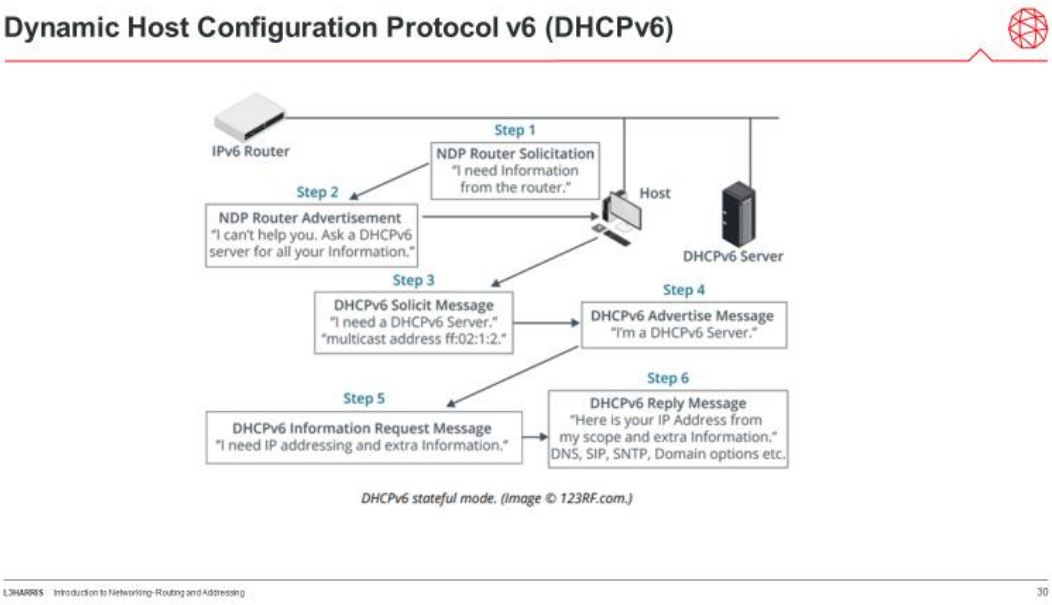

IPv6 can locate routers (default gateways) and generate a host address with a suitable network prefix automatically as part of Stateless Address Autoconfiguration (SLAAC). In this context, the role of a DHCP server in IPv6 is different. In an IPv6 network, DHCPv6 is often just used to provide additional option settings, rather than leases for host IP addresses. DHCPv6 is defined in RFC 3315. The format of messages is different, but the process of DHCP server discovery and address leasing (if offered) is fundamentally the same.

DHCPv6 Stateless: In stateless mode, a client obtains a network prefix from a Router Advertisement and uses it with the appropriate interface ID. The router can also set a combination of flags to tell the client that a DHCP server is available.

DHCP Stateful and Prefix delegation (PD): Stateful mode means that a host can also obtain a routable IP address from a DHCPv6 scope, plus any other options similar to DHCP for IPv4).

- Configuring the scope requires you to define the network prefix and then any IP addresses that are to be excluded from being offered.
- All other addresses that are not explicitly excluded can be offered.
- The host must still listen for a router advertisement to obtain the network prefix and configure a default gateway.
- There is no mechanism in DHCPv6 for setting the default route.

DHCPv6 Prefix Delegation (PD) is used by ISPs to provide routable address prefixes to a SOHO router, installed as customer premises equipment (CPE). With PD, the CPE router obtains a prefix from a delegating router, installed upstream on the ISP's network. The CPE router then uses the prefix to assign devices on the customer's network with IPv6 addressing information by using router advertisements or DHCPv6, or both.

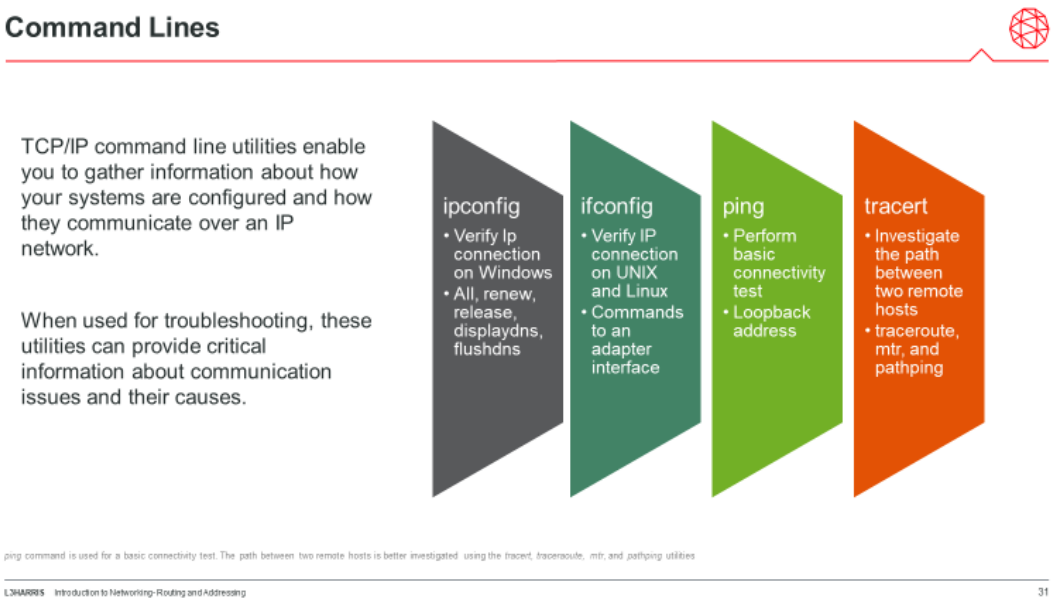

***ipconfig* command;** is used to verify the IP configuration on Windows-based systems and includes the following options:

- *ipconfig* without any switches will display the IP address, subnet mask, and default gateway (router) for all network interfaces to which TCP/IP is bound.
- *ipconfig /all* displays complete TCP/IP configuration parameters for each interface to which TCP/IP is bound, including whether the Dynamic Host Configuration Protocol (DHCP) is enabled for the interface and the interface's hardware (MAC) address.
- *ipconfig /renew* interface forces a DHCP client to renew the lease it has for an IP address.
- *ipconfig /release* interface releases the IP address obtained from a DHCP Server so that the interface(s) will no longer have an IP address.
- *ipconfig /displaydns* displays the Domain Name System (DNS) resolver cache.
- *ipconfig /flushdns* clears the DNS resolver cache. • ipconfig /registerdns registers the host with a DNS server (if it supports dynamic updates).

ifconfig: UNIX® and Linux® hosts provide the ifconfig command, ifconfig, which provides similar output to the Windows® ipconfig program.

- ifconfig can also be used to bind an address to an adapter interface, set up communications parameters, and enable or disable the adapter.

***ping*:** The ping utility sends a configurable number and size of ICMP packets to a destination host.

***tracert**:* trace the route taken by a packet as it hops to the destination host on a remote network. It can be used either with an IP address or a host and domain name. It returns the IP address (or name) of each router used by the packet to reach its destination.

ping Command

Verfiy a host's configuration and test for router connections:

1. Ping the loopback address (ping l27.0.0.l) to verify TCP/IP is installed and loaded correctly.
2. Ping the IP address of your workstation to verify it was added correctly and to verify that the network adapter is functioning properly.
3. Ping the IP address of the default gateway to verify it is up and running and that you can communicate with a host on the local network.
4. Ping the IP address of a remote host to verify you can communicate through the router.

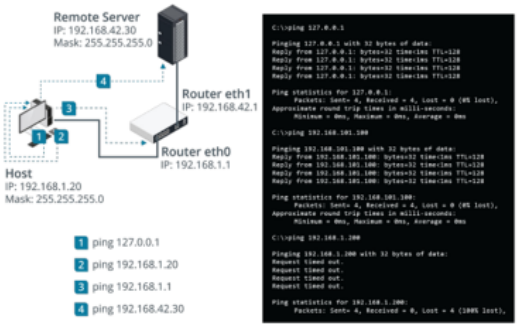

Troubleshooting with ping. (Image © 123RF.com. Screenshot used with permission from Microsoft.)

***ping*:** The ping utility sends a configurable number and size of ICMP packets to a destination host. This can be used to perform a basic connectivity test that is not dependent on the target host running any higher-level applications or services.

To use *ping,* open a command prompt and enter this command: *ping HostName* or *IPAddress* in reference to the remote computer.

- **Successful:** responds with the message "Reply from *IPAddress*" and the time it takes for the server's response to arrive. The millisecond measures of Round Trip Time (RTT) can be used to diagnose latency problems on a link
- **Destination host unreachable**—There is no routing information (that is, the local computer does not know how to get to that IP address). If the host is on the same IP network, check physical cabling, infrastructure devices such as the switch, and IP configuration. If the host is on another IP network, check the IP configuration and router.
- **No reply** (Request timed out.)—The host is unavailable or cannot route a reply to your computer.

ping Output Analysis

The trick with ping is understanding the messages that you receive when there is a problem.

- If you cannot ping the loopback address, the protocol is not correctly installed on the local system.
- If you cannot ping your own address, there might have been a configuration error, or the network adapter or adapter driver could be faulty.
- If a local host cannot be pinged, then verify the sending host's IP configuration—IP address, subnet mask, and so on.
- If the previous tests are successful, but a remote IP address cannot be contacted, check the default gateway parameter on the local host. If correct, use the tracert command, tracert, to investigate the route being taken. Also, consider manually adding the route by using the route command.
- If you can successfully perform all tests by IP address, but cannot ping by computer name, then this suggests a name resolution problem

ping SWITCHES

ping can be used with several switches.

- In Windows, you can adjust the TTL (-i) and timeout (-w) and force the use of IPv4 (-4) or IPv6 (-6) when pinging by host name.
- With IPv4, you can also use loose (-j) or strict (-k) source routing (sending packets via a predetermined route).
- The -a switch performs name resolution.
- Also, -t continues to ping the host until interrupted (by pressing Ctrl+C).

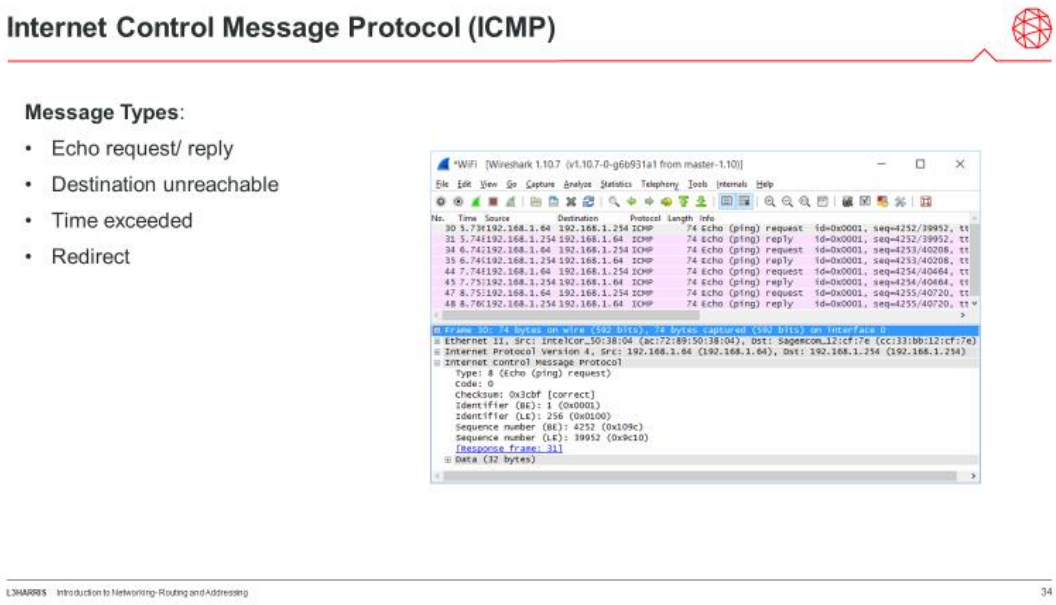

Internet Control Message Protocol (ICMP): is used to report errors and send messages about the delivery of a packet. It can also be used to test and troubleshoot connectivity issues on IP networks. ICMP messages are generated under error conditions in most types of unicast traffic, but not for broadcast or multicast packets. An ICMP message is encapsulated within a single IP packet. ICMP messages are categorized into various types, the most commonly encountered are as follows;

- **Echo Request/ Reply:** These are used for testing a connection with the ping utility. If a request message reaches the destination host, it generates a reply and sends it back to the source. If the request message does not reach its destination, an appropriate error message is generated.
- **Destination Unreachable**: This class of message indicates that a local host or a host on a remote network (or a protocol or port on a host) cannot be contacted. This might be caused by some sort of configuration error or by a host or router not being available.
- **Time Exceeded:** This is used when the Time to Live (TTL) of a packet reaches 0. The TTL field in a packet has a maximum value of 255, and this value is reduced by one every time the packet crosses a router. The TTL is eventually reduced to 0 if the packet is looping (because of a corrupted routing table) or when congestion causes considerable delays. The router then discards the packet, and a warning packet is sent back to the source host.
- **Redirect**: Most hosts channel all remote communications through the default gateway. If there are in fact multiple routers and a more efficient route can be identified, the default gateway can send a redirect message to the host to update its routing table. The router still delivers the original message.

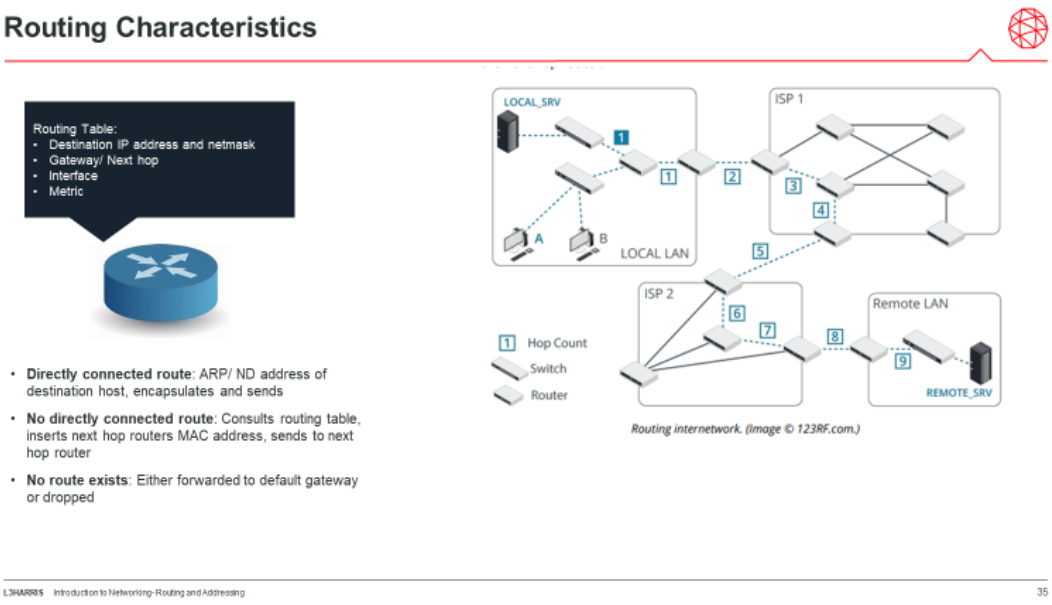

- **End systems (ESs):** Hosts with no capacity to forward packets to other IP networks
- **Intermediate Systems (ISs):** Routers that interconnect IP networks and can perform the packet forwarding process.

Routing tables: On a router, information about the location of other IP networks and hosts is stored in a routing table. Each entry in the routing table represents an available route to a destination network or host, and contains (at least) the following parameters:

- Destination IP address and netmask: Routes can be defined to specific hosts but are more generally directed to network IDs.
- Gateway/next hop: The IP address of the next router along the path.
- Interface: The local port to use to forward a packet along the chosen route.
- Metric: A preference value assigned to the route, with low values being preferred over high ones. The value of the metric may be determined by different parameters. (ie. how far the next hop router is, how long it will take to route a packet to the subsequent routers, available bandwidth, how large a packet can be sent without fragmentation, etc.)

When a router receives a packet, it goes through the same process that the source host did to calculate if the packet needs to be routed to another router or it can be delivered locally to another interface.

If the packet has been routed, the Time to Live (TTL) is decreased by at least 1. When the TTL is 0, the packet is discarded.

- **Static routing** means that you manually add routes to a routing table, and they change only if you edit them
- A **default route** is a special type of static route that identifies the next hop router for an unknown destination. (0.0.0.0/0 (IPv4) or ::/0 (IPv6))

Dynamic Routing Protocols

Routers use protocols to exchange information about connected networks and select the best available route to a destination. These protocols use various algorithms and metrics to build and maintain routing tables to provide reasonably current routing information about the networks to which they are connected.

Protocol	Type	Class	Transport
Routing Information Protocol (RIP)	Distance- vector	IGP	UDP (port 520 or 521)
Enhanced Interior Gateway Routing (EIGRP)	Distance-vector hybrid	IGP	Native IP (88)
Open Shortest Path First (OSPF)	Link-state	IGP	Native IP (88)
Border Gateway Protocol (BGP)	Distance-vector hybrid	EGP	TCP (port 179)

Routing Information Protocol (RIP) uses a hop count metric to determine the distance to the destination network. To help prevent looping, the maximum hop count allowed is 15. Consequently, this limits the maximum size of a RIP network, since networks that have a hop count of 16 or higher are unreachable. **RIPv2** supports classless addressing. **RIPng** supports IPv6

Enhanced Interior Gateway Routing Protocol (EIGRP) uses a metric composed of several administrator weighted elements and supports multiple paths to the destination network.

- A native IP protocol, which means that it is encapsulated directly in IP datagrams and tagged with the protocol number 88 in the protocol field of the IP header.

Open Shortest Path First (OSPF): is a hierarchical link-state routing protocol that was designed to support classless addressing.

- Networks and their connected hosts and routers within an autonomous system are grouped into OSPF areas.
- Routers within a given area share the same topological database of the networks they serve. Routers that can connect to multiple areas are known as area border routers
- Routers use a Link State Advertisement (LSA) to update their routing tables. These exchanges of routing information enable the routers to each build a topological routing tree and keep it up to date.

Border Gateway Protocol (BPG): is designed to be used between routing domains, or Autonomous Systems (ASes), and as such is used as the routing protocol on the Internet, primarily between ISPs

- Border (or edge) routers for each AS exchange only as much network reachability information as is required to access other autonomous systems (the AS path)
- Autonomous System Numbers (ASN) are allocated to ISPs by IANA via regional registries.

Router Configuration

- A router can be implemented as hardware or software.
- Routers often support the functions of a firewall.
- WAN router provides access to the Internet.
- Basic switches provide ports and virtual LANs for wired and wireless (via an access point) devices.
- Traffic between logical networks is controlled by a LAN router (or layer 3 switch).
- A router may be a dedicated appliance with a port to each of the networks, or it may be a network operating system (NOS) server with multiple interface cards.

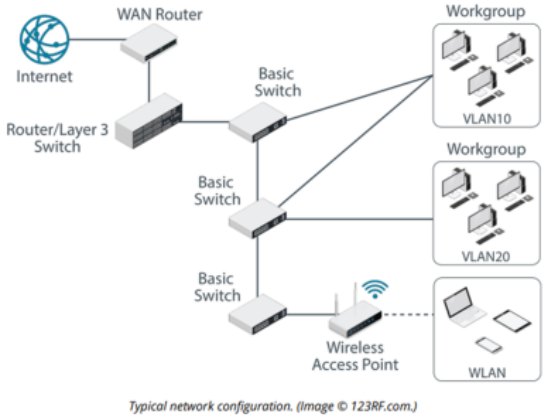

Typical network configuration. (Image © 123RF.com.)

Routers serve both to join physically remote networks and subdivide autonomous IP networks into multiple subnets. Border or edge routers are typified by distinguishing external (Internet-facing) and internal interfaces. These devices are placed at the network perimeter.

Hot Standby Router Protocol (HSRP) developed by Cisco in 1998, HSRP allows for multiple physical routers to serve as a single default gateway for a subnet.

- Each router must have an interface connected to the subnet, with its own unique MAC address and IP address as well as be configured to share a common virtual IP address and a common MAC address.
- They communicate among themselves using IP multicasts and choose an active router based on priorities configured by an administrator the reaming routers form a standby group, the router with the next highest priority is chosen as the standby router and will take over if the active router becomes unavailable.

Virtual Router Redundancy Protocol (VRRP) Developed in 2004, the current version is VRRP version 3 which adds support for IPv6.

- In VRRP, the active router is known as the master, and all other routers in the group are known as backup routers
- It is possible to configure VRRP routers to use only the virtual IP address.

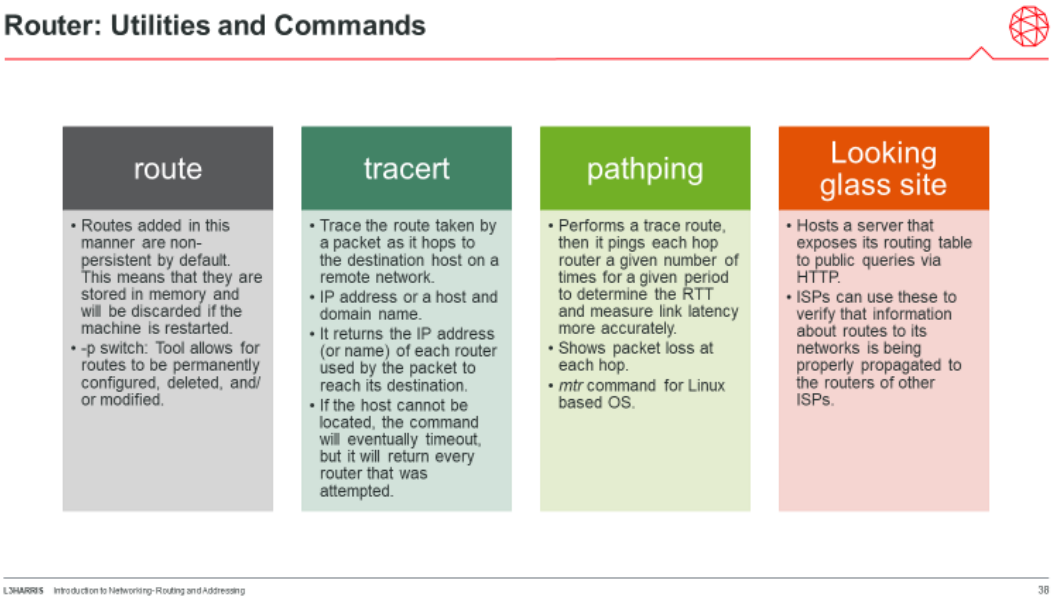

A host's routing table contains information about routes to other hosts. A router for a complex network would normally have a large routing table populated dynamically by one or more routing protocols. An end system will usually have a simple routing table configured with a few default entries. For example, the default entries for a Windows® host are:

- Default route (0.0.0.0/0).
- Loopback address.
- Host's subnet address.
- Host's own address.
- Multicast address.
- Broadcast address.

Transmission Control Protocol

- Orderly connection establishment and teardown
- Segmentation
- Acknowledgements
- Flow Control

Field	Explanation
Source Port	TCP port of sending host
Dest. Port	TCP Port of dest. host
Sequence #	ID # of current segment
ACK #	Sequence # of next segment expected
Data Length	Length of TCP segment
Flags	Type of content in the segment
Window	Amount of data host is willing to receive before sending another acknowledgement
Checksum	Ensures validity of the segment
Urgent Pointer	Specifies the end of urgent data in a segment
Options	Further configures , ie Maximum Segment Size

In the TCP/IP suite, the Transmission Control Protocol (TCP) provides a connection oriented, guaranteed method of communication using acknowledgements to ensure delivery. TCP takes data from the Application layer as a stream of bytes and divides it up into segments, each of which is given a header. The TCP segments become the payload of the underlying IP datagrams. TCP requires that a connection be established before hosts can exchange data. A connection uses the following features to ensure reliability

Transmission Control Protocol (TCP): provides a connection oriented, guaranteed method of communication using acknowledgements to ensure delivery.

- **Orderly connection establishment and teardown**—The client and server perform a handshake to establish and end connections. Under normal circumstances, a single connection is created between hosts. However, multiple connections can be established by a single application process to improve throughput.
- **Segmentation**—TCP breaks PDUs from the Application layer into a segment format and uses sequence numbers to allow the receiver to rebuild the message correctly. This allows the connection to deal with out-of-order packets.
- **Acknowledgements** (ACKs)—Packets might be out-of-order because they are delayed, but they could also be lost completely or arrive in a damaged state. In the first case, the lack of acknowledgement results in the retransmission of 0 data and, in the second case, a Negative Acknowledgement (NAK or NACK) forces retransmission.
- **Flow control**—Enables one side to tell the other when the sending rate must be slowed.

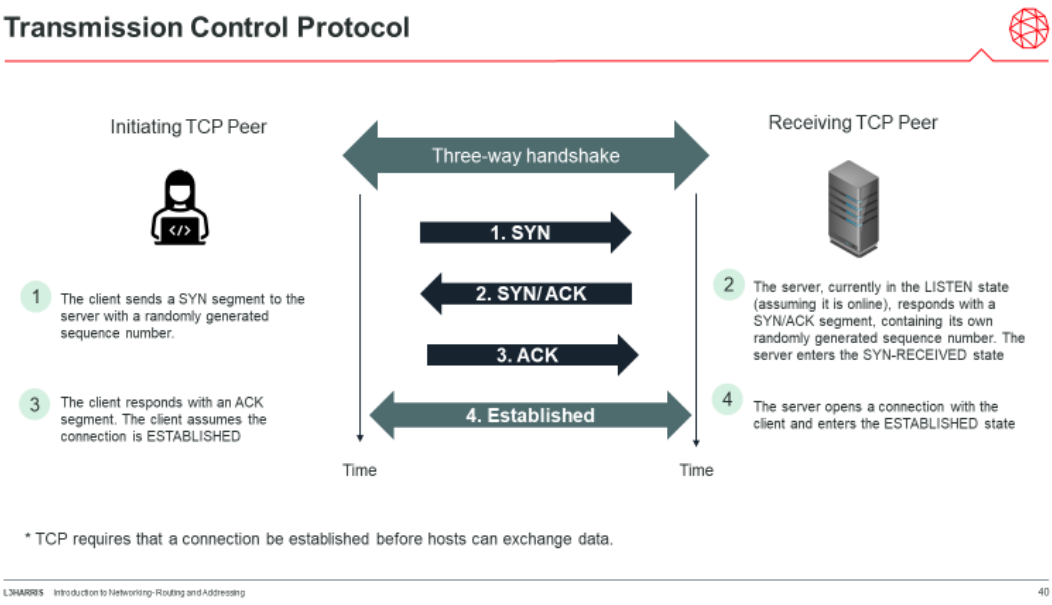

TCP Three-way handshake and Flow control: A TCP connection is typically established to transfer a single file, so a client session for something like a web page (HTTP) might involve multiple TCP connections being opened with the server. A connection is established using a three-way handshake

The sending machine expects regular acknowledgements for segments it sends and, if a period elapses without an acknowledgement, it assumes the information did not arrive and automatically resends it. This overhead makes the system relatively slow. Connection-oriented transmission is suitable when reliability and data integrity are important. Another important function of TCP is handling flow control to make sure the sender does not inundate the receiver with packets. The main mechanism used for this is called the sliding window. The window field in the header represents the number of bytes starting from the last acknowledged byte that a host is prepared to receive before it will send an acknowledgement.

There are also functions for resetting a connection and (in some implementations) keeping a connection alive if no actual data is being transmitted (hosts are configured to timeout unused connections).

User Datagram Protocol (UDP)

- Connectionless, non-guaranteed method of communication with no sequencing or flow control.
- Used by Application layer protocols that need to send multicast or broadcast traffic (TCP supports unicast only).
- Used for applications that transfer time-sensitive data but do not require complete reliability, such as voice or video

Port #	Service/ Application	Description
53*	Domain	Domain Name System
67	bootps	BOOTP/DHCP/ Server
68	bootpc	BOOTP/DHCP Client
69	tftp	Trivial file transfer protocol
123	ntp	Network Time Protocol
161	snmp	Simple Network Management Protocol
162	Snmp-trap	Simple Network Management Protocol Trap
546	dhcpv6-client	DHCPv6 Client
547	dhcpv6-server	DHCPv6 Server
5004	rtp	Real-Time Protocol
5005	rtcp	Real-Time Control Protocol
5060*	sip	Session Initiation Protocol
5061*	sips	Session Initiation Protocol Secure

* TCP/UDP

The **User Datagram Protocol (UDP)** also works at the Transport layer, but unlike TCP, it is a connectionless, non-guaranteed method of communication with no sequencing or flow control. There is no guarantee regarding the delivery of messages or the sequence in which packets are received. The reduced overhead means that delivery is faster.

Structure of a UDP datagram:

- Source port- UDP port of sending host
- Destination Port- UDP port of destination host
- Message Length- The size of this UDP message
- Checksum- Verify the datagram

The header size is 8 bytes, compared to 20 bytes (or more) for TCP

TCP and UDP Ports: Any application or process that uses TCP or UDP for its transport is assigned a unique identification number called a port. Ports are logically assigned to provide a communications channel between a server application on one host and a client application on another host, allowing them to send and receive data. Port numbers for some server application protocols are pre-assigned by the Internet Assigned Numbers Authority (IANA). IANA assigns protocols to port numbers 0 through 1023. These port assignments are documented at iana.org/assignments/servicenames-port-numbers/service-names-port-numbers.xhtml. Vendors can register ports 1024 through 49,151.

The port number is used in conjunction with an IP address to form a socket. A socket provides an endpoint to a connection, and two sockets form a complete path. A socket works as a bi-directional pipe for incoming and outgoing data.

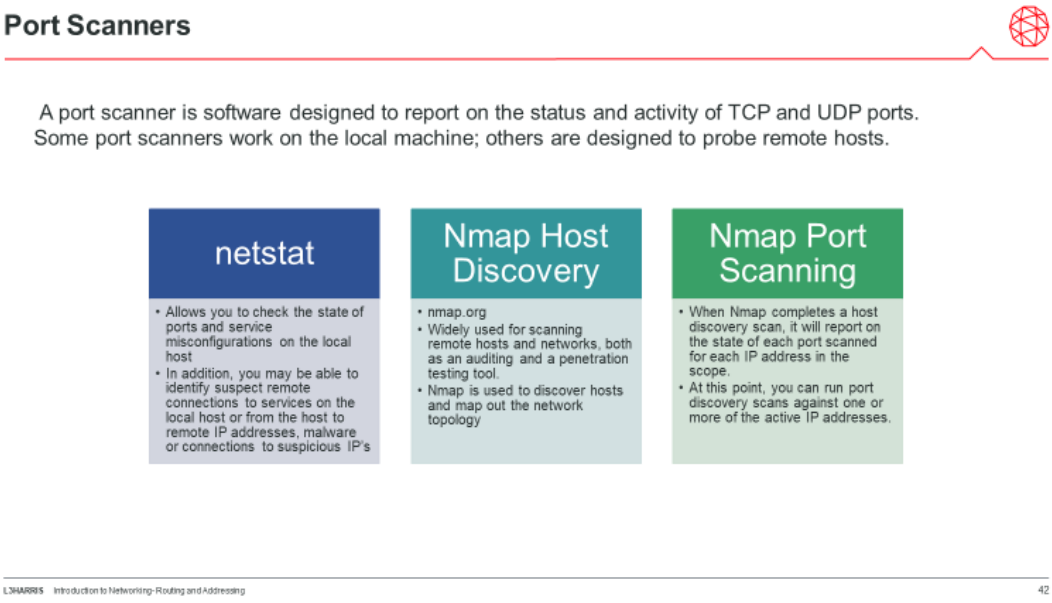

Netstat: On Windows® , used without switches, the command outputs active TCP connections, showing the local and foreign addresses and ports.

The following additional switches can be used:

- -a displays all connections (active TCP and UDP connections plus ports in the listening state).
- -o shows the Process ID (PID) number that has opened the port.
- -b shows the process name that has opened the port.
- -n displays ports and addresses in numerical format. Skipping name resolution speeds up each query.
- -s shows per protocol statistics (such as packets received, errors, discards, unknown requests, port requests, failed connections, and so on).
- -p proto displays connections by protocol (TCP or UDP or TCPv6/UDPv6). When used with -s, this switch can also filter the statistics shown by IP, IPv6, ICMP, and ICMPv6.

Nmap Host Discovery: The Nmap Security Scanner (nmap.org) is widely used for scanning remote hosts and networks, both as an auditing and a penetration testing tool.

- As well as port scanning, Nmap is used to discover hosts and map out the network topology.
- Nmap can use diverse methods of host discovery, some of which can operate stealthily and serve to defeat security mechanisms such as firewalls and intrusion detection.

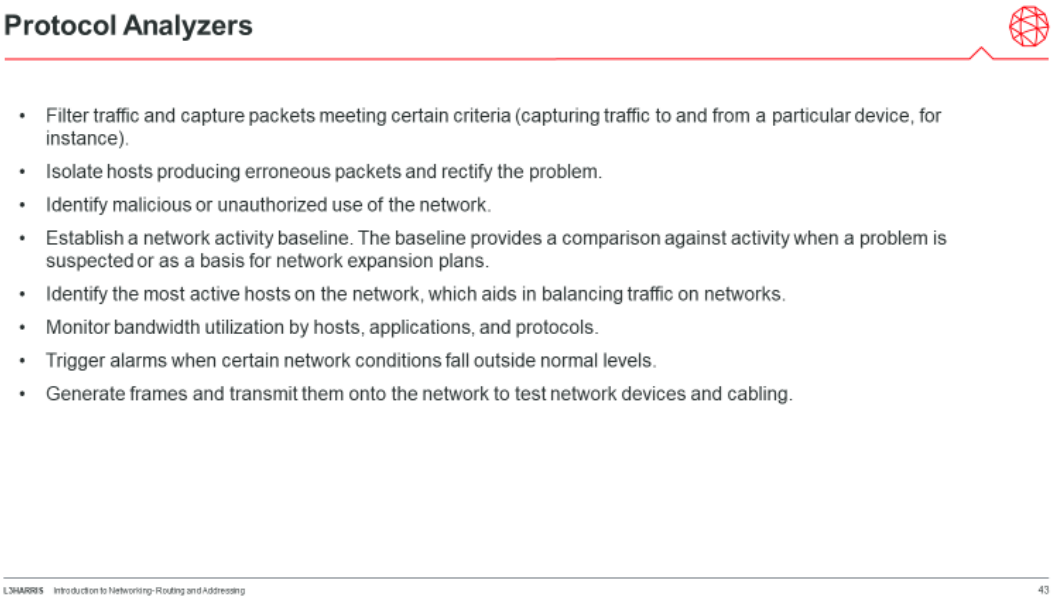

Protocol Analyzer: A protocol analyzer (or packet analyzer) works in conjunction with a packet sniffer. You can either analyze a live capture or open a saved capture (.pcap) file. Protocol analyzers can decode a captured frame to reveal its contents in a readable format. You can choose to view a summary of the frame or choose a more detailed view that provides information on the OSI layer, protocol, function, and data. Analyzing protocol data at the packet level will help to identify protocol or service misconfigurations. You can also perform traffic analysis to monitor statistics related to communications flows, such as bandwidth consumed by each protocol or each host, identifying the most active network hosts, monitoring link utilization and reliability, and so on.

Wireshark: is an open source graphical packet capture and analysis utility, with installer packages for most operating systems**.** Wireshark is capable of parsing (interpreting) the headers of hundreds of network protocols. You can apply a capture filter using the same expression syntax as tcpdump. You can also apply display filters by using a different and more powerful set of expressions (a query can be built via the GUI tools, too). Another useful option is to use the Follow TCP Stream context command to reconstruct the packet contents for a TCP session. Use the Statistics menu to access traffic analysis tools.

Domain Name System (DNS)

- DNS is operated by ICANN (icann.org), which also manages the generic TLDs..
- Each domain name must be registered with a Domain Name Registry for the appropriate top-level domain..
- Information about a domain is found by tracing records from the root down through the hierarchy.
- The root servers have complete information about the top-level domain servers. In turn, these servers have information relating to servers for the second level domains.
- No name server has complete information about all domains.

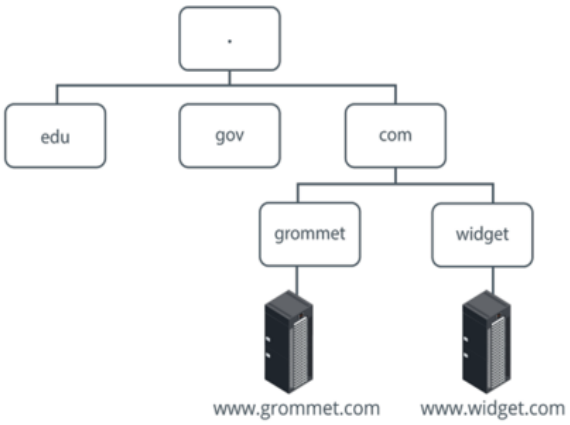

DNS hierarchy. (Image © 123RF.com.)

To avoid the possibility of duplicate host names on the Internet, a **fully qualified domain name (FQDN)** is used to provide a unique identity for the host belonging to a particular network.

- An FQDN is made up of the host name and a domain suffix.
- Numerous hosts may exist within a single domain.
- FQDNs must follow certain rules

Domain Name System (DNS): is a hierarchical system of distributed name server databases that contain information on domains and hosts within those domains. The system's distributed nature has the twin advantages that maintenance is delegated and loss of one DNS server does not necessarily prevent name resolution from being performed. At the top of the DNS hierarchy is the root, which is represented by the null label, consisting of just a period (.). There are 13 root level servers (A to M) Immediately below the root lie the top-level domains (TLDs). There are several types of top-level domains, but the most prevalent are generic (.com, .org, .net, .info, .biz), sponsored (.gov, .edu), and country code (.uk, .ca, .de).

The signal for the name resolution process to commence occurs when a user presents an FQDN (often within a URL) to an application program, such as a web browser. The client application, referred to as a stub resolver, checks its local cache for the mapping. If no mapping is found, it forwards the query to its local name server. The IP addresses of primary and secondary (backup) name servers are usually set in the TCP/IP configuration.

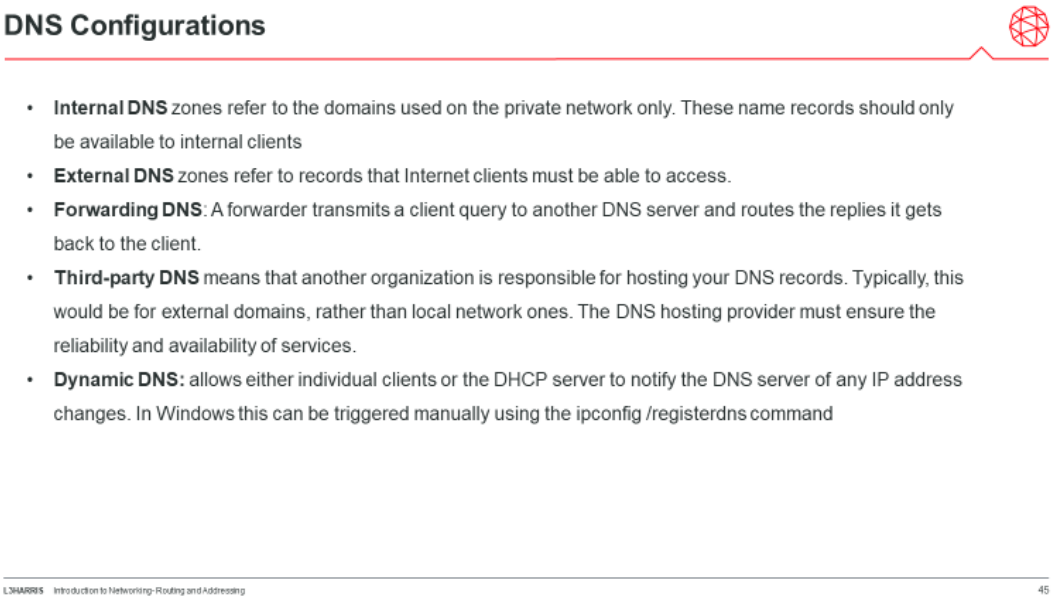

A name server can maintain primary and secondary zones: **Primary** means that the zone can be edited. **Secondary** means a read-only copy of the zone.

Resource Records: A DNS zone will contain numerous resource records. These records allow the DNS server to resolve queries for names and services hosted in the domain into IP addresses. Resource records can be created and updated manually (statically) or dynamically from information received from client and server computers on the network. Multiple different named resource records can refer to the same IP address. The most common records are:

- **Start of Authority (SOA)**: identifies the primary DNS name server that is authoritative for the zone and is therefore responsible for resolving names in the domain (plus any subdomains). The SOA also includes contact information for the zone and a serial number (for version control).
- **Name Server (NS**): identify authoritative DNS name servers for the zone. In most enterprise networks, each zone will have several (at least two) DNS servers holding a replicated copy of the zone. Therefore, two or more NS records are configured for redundancy.
- **A Record**: is used to resolve a host name to an IPv4 address.
- **Canonical Name (CNAME)** record is used to represent an alias for a host.
- **Mail Exchanger (MX) record** is used to identify an email server for the domain.
- **TXT record** is used to store any free-form text that may be needed to support other network services. A single domain name may have many TXT records, but most commonly they are used as part of Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM)
- **Pointer Record** :A DNS server may have two types of zones: forward lookup and reverse lookup. Forward lookup zones contain most of the resource records you are looking at here— Given a name record, a forward lookup returns an IP address. A PTR record is found in reverse lookup zones and is used to resolve an IP address to a host name.

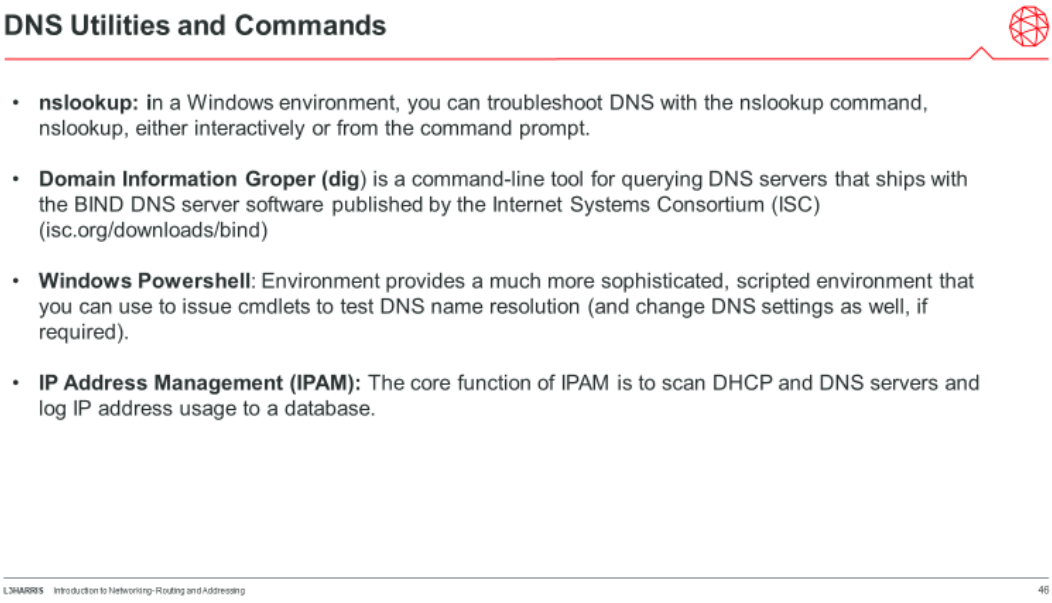

nslookup: in a Windows environment, you can troubleshoot DNS with the nslookup command, nslookup, either interactively or from the command prompt.

nslookup -Option Host Server :Host can be either a host name, a domain name, a Fully Qualified Domain Name (FQDN), or an IP address. Server is the DNS server to query; the default DNS server is used if this argument is omitted.

- If nslookup is run without any arguments (or with just the argument -Server), the tool is started in interactive mode. You can perform specific query types and output the result to a text file for analysis

IP Address Management: The core function of IPAM is to scan DHCP and DNS servers and log IP address usage to a database. Most suites can scan IP address ranges to detect use of statically assigned addresses. Some IPAM software may also be able to scan the hardware associated with an IP address and save the information to an asset inventory. IPAM software can often be used to manage and reconfigure DHCP and DNS servers remotely. The software also provides analysis tools to allow administrators to identify overloaded DHCP scopes or to make more valuable public IP addresses available. IPAM also performs valuable incident response and forensics functions. IPAM may reveal unauthorized use of IP addresses or address ranges and can track the use of an IP address over time.

- Windows Server ships with a basic IPAM tool, though it is only suitable for managing Windows-based servers. Cisco has their Prime Network Registrar (cisco.com).
- Other popular IPAM vendors include Infoblox (infoblox.com) and Efficient IP (efficientip.com)

Domain Information Groper (dig) is a command-line tool for querying DNS servers that ships with the BIND DNS server software published by the Internet Systems Consortium (ISC) (isc.org/downloads/bind)

Windows Powershell: Environment provides a much more sophisticated, scripted environment that you can use to issue cmdlets to test DNS name resolution (and change DNS settings as well, if required).

Summary Utilities and DHCP

- ipconfig/ifconfig is used for various troubleshooting and configuration tasks—There are differences between the Windows and Linux versions.
- ICMP delivers status message and allows for connectivity testing (ping utility).
- DHCP is a method for a client to automatically request IP configuration information from a server.
- A DHCP service can be configured to run on a Windows/Linux server, or it can be provided by most types of switches and routers.
- In Windows, if a client cannot contact a DHCP server, it uses an APIPA (link-local) address that allows for communications on the local network.
- DHCPv6 is principally used to provide optional information, such as DNS server addresses. Router advertisements can be used to direct a client to use a stateless or stateful autoconfiguration