Short answer

What is the purpose of the Internet Engineering Task Force (IETF)? is the group responsible for drafting, testing, proposing, and maintaining official Internet Standards, in the form of RFCs, through the agencies of multiple working groups under its purview.

The reference model described in ISO Standard 7498 breaks network communication into seven layers. List each layer from top to bottom.

Application layer

Presentation layer

Session layer

Transport layer

Network layer

Data Link layer

Physical layer

Provide brief descriptions of the following protocols: High-level Data Link Control (HDLC) protocol and frame relay. High-level Data Link Control (HDLC) protocol: Based on IBM's original SNA Data Link Control (SDLC) protocol. HDLC uses data frames to manage network links and data transmission.

Frame relay: A telecommunications service designed to support intermittent data transmission between local area networks and wide area network end points. Frame relay uses data frames to manage network links and data transmission.

Briefly describe the three primary tasks that the Internet layer handles for TCP/IP. MTU fragmentation: When a route carries data from one type of network to another, the largest chunk of data that the network can carry, an

MTU, can vary. When data moves from a medium that supports a larger MTU to a medium that supports a smaller MTU, that data must be reduced to smaller pieces to match the smaller of the two MTUs involved.

Addressing: This defines the mechanism whereby all network interfaces on a TCP/IP network must be associated with specific, unique bit patterns that identify each interface individually, and also identify the network (or even network locale) to which that interface belongs.

Routing: This defines the mechanism that forwards packets from sender to receiver, in which numerous intermediate relays may be involved in achieving delivery from sender to receiver.

What is the purpose of the following protocols: Internet Protocol, Internet Control Message Protocol, and Address Resolution Protocol. Internet Protocol (IP): Routes packets from sender to receiver.

Internet Control Message Protocol (ICMP): Handles information about IP-based routing and network behavior, especially as they relate to " traffic conditions" and errors.

Address Resolution Protocol (ARP): Address Resolution Protocol (ARP) converts between numeric IP network addresses and Media Access Control (MAC) addresses on a specific cable segment (always used for the final step of packet delivery).

Routing: This defines the mechanism that forwards packets from sender to receiver, in which numerous intermediate relays may be involved in achieving delivery from sender to receiver.

What is the difference between the Open Shortest Path First protocol and the Border Gateway Protocol? Open Shortest Path First (OSPF): Defines a widely used, link-state routing protocol for local or interior routing regions within local internetworks.

Border Gateway Protocol (BGP): Defines a widely used routing protocol that connects to common Internet backbones, or other routing domains within the Internet where multiple parties jointly share responsibility for managing traffic.

Briefly discuss two elements that TCP/IP services depend on to operate. In UNIX terminology, a special " listener process," called a daemon, operates on a server to handle incoming user requests for specific services. On Windows Server 2008, a process called INETINFO. EXE appears in the Task Manager's Processes tab whenever the Web server, IIS, or FTP server is running.

Each TCP/IP service has an associated port address that uses a 16-bit number to identify a specific process or service. Addresses in the range from 0 to 1024 are often called well-known port addresses and associate a specific port address with a specific service.

Briefly describe three options for analyzing switched networks. Hubbing out: By placing a hub between a device of interest (such as a server) and the switch, and connecting the analyzer to the hub, you can view all traffic to and from the server.

Port redirection: Many switches can be configured to redirect (actually, to copy) the packets traveling through one port to another port. By placing your

analyzer on the destination port, you can listen in on all the conversations that cross the network through the port of interest.

Remote Monitoring (RMON): Uses Simple Network Management Protocol (SNMP) to collect traffic data at a remote switch and send the data to a management device.

Briefly discuss IPs three-part addressing scheme. Symbolic: This consists of names that take a particular form, such as www. support. dell. com.

Logical numeric: This consists of a set of four numbers, separated by periods, as in 172. 16. 1. 10. Each of these four numbers must be smaller than 256 in decimal to be represented in eight binary digits, or bits.

Physical numeric: This consists of a six-byte numeric address, burned into firmware (on a chip) by network interface manufacturers.

Why are concepts such as subnets and supernets important for TCP/IP networks? each of these ideas refers to a single " local neighborhood" on such a network, seen from a routing perspective. When network addresses are further subdivided beyond their defaults for whatever class to which an address belongs, such subnetting represents " stealing bits" (borrowing bits) from the host portion of the address and using those stolen (borrowed) bits to create multiple routing regions within the context of a single network address.

Briefly describe how to calculate subnet masks. The simplest form of subnet masking uses a technique called constant-length subnet masking (CLSM), in which each subnet includes the same number of stations and represents a

simple division of the address space made available by subnetting into multiple equal segments.

Another form of subnet masking uses a technique called variable-length subnet masking (VLSM) and permits a single address to be subdivided into multiple subnets, in which subnets need not all be the same size.

What are the limitations of creating a CIDR address? 1. All the addresses in the CIDR address must be contiguous. Use of the standard network prefix notation for addresses, however, also makes it tidy and efficient to carve up any kind of address, as needed.

2. When address aggregation occurs, CIDR address blocks work best when they come in sets that are greater than 1 and equal to some lower-order bit pattern that corresponds to all 1s - namely in groups of 3, 7, 15, 31, and so on. That's because this makes it possible to borrow the corresponding number of bits (two, three, four, five, and so on) from the network portion of the CIDR address block and use them to extend the host portion instead.

3. To use a CIDR address on any network, all routers in the routing domain must " understand" CIDR notation. This is typically not a problem for most routers that were built after September 1993, when RFCs 1517, 1518, and 1519 were approved, because most router vendors began to support CIDR addresses at that time.

What are the disadvantages of using private IP addresses? Such addresses may not be routed across the public Internet.

Some IP services require what's called a secure end-to-end connection - IP traffic must be able to move in encrypted form between the sender and receiver without intermediate translation. Thus, if either party to such a connection uses a public IP address, it's easiest to configure if both parties use a public IP address because the address for the " private end" of the connection cannot be routed directly across the Internet.

Most organizations need public IP addresses only for two classes of equipment. Briefly describe each of these classes. Devices that permit organizations to attach networks to the Internet. These include the external interfaces on boundary devices of all kinds, such as routers, proxy servers, and firewalls, that help maintain the perimeter between the " outside" and " inside" on networks.

Servers that are designed to be accessible to the Internet. These include public Web servers, e-mail servers, FTP servers, news servers, and whatever other kind of TCP/IP Application layer services an organization may want to expose on the public Internet.

List the constraints that determine the number and size of networks. Number of physical locations
Number of network devices at each location
Amount of broadcast traffic at each location
Availability of IP addresses
Delay caused by routing from one network to another

Give two reasons why you should use binary boundaries. One reason is that, in the future, you may want to implement layer-3 switching to reduce the

broadcast traffic, and if the devices fit in a binary boundary, you won't have to readdress them.

Another good reason to use binary boundaries is that one day you will want to classify your traffic to apply Quality of Service (QoS) or policies of some sort.

What are some of the design goals for IPv6? Pv6 not only provides a vast abundance of IP addresses and better management of its address space, it eliminates the need for NAT and other technologies to be put in place to shore up the inadequate number of IPv4 addresses. IPv6 also makes it easier to administer and configure IP addresses. Also, IPv6 has modernized routing support and natively allows for expansion
along with the growing Internet.
Finally, IPv6 supports network security by using authentication and encryption extension headers, among other methods.

Describe the IPv4 Header Length field. The Header Length field, also referred to as the Internet Header Length (IHL) field, denotes the length of the IP header only the IP header can support options and, therefore, may vary in length. IHL includes an offset to the data to make it fall on a 32-bit boundary value. The minimum value for IHL is 5, this produces a length of 5 * 32 = 160 bits, which equals 20 bytes. The IHL is a 4-bit field, so the maximum value it can represent is 16 - 1 = 15.

What is a real-time application (RTA)? Real-time applications (RTAs), in the current context, are any applications that must function within an immediate time frame on a continual basis, with little or no latency (delay). Whether a

service can be defined as RTA depends on the maximum amount of time an application task requires to be executed on a particular hardware platform. This is referred to as worst-case execution time (WCET).

Voice over IP (VoIP), which uses IP to support voice communication, is one type of RTA

Is packet fragmentation in IPv4 a good thing? Sometimes. A packet can only successfully arrive at its destination if the MTU is supported by the smallest link MTU in the path. IPv4 routers continually fragment traffic from one hop to the next depending on the MTU size tolerated by the next link, and in IPv6, source nodes use PMTU Discovery to determine the smallest link MTU in a path and then set the MTU of its packets to accommodate that link.

What is the role of the Next Header field in IPv6? The 8-bit Next Header field specifies the header type of the header immediately following the IPv6 header —specifically, extension headers —and uses the same values as the IPv4 Protocol field.

What is the recommended extension header ordering in IPv6? 1. Hop-by-Hop Options

2. Destination Options

3. Routing

4. Fragment

5. Authentication

6. Encapsulating Security Payload (ESP)

Briefly describe the IPv6 Authentication extension header. The Authentication extension header is designed to specify the true origin of a packet by preventing address spoofing and connection theft. This header also provides an integrity check on those parts of the packet that do not change in transit. (Authentication would not be calculated over the Routing extension header, for example.) In addition, the Authentication extension header can provide a limited defense against replay attacks. End devices may, if configured to do so, reject packets that are not properly authenticated.

The Authentication extension header starts with a 1-byte Next Header field that indicates the next header in the chain

What is " 6to4"? Connection of IPv6 Domains via IPv4 Clouds," specifies an optional method for IPv6 sites to communicate with one another over IPv4 networks without setting up tunneling. This mechanism treats IPv4 wide area networks (WANs) as a unicast point-to-point link. It is a temporary fix

What are the two most important jobs of the Data Link layer? Managing access to whatever networking medium is in use, called Media Access Control (usually abbreviated as MAC)

Creating temporary point-to-point links between a pair of MAC layer addresses to enable data transfer, called Logical Link Control (usually abbreviated as LLC)

Briefly describe the following fields in the PPP header and trailer: Flag, protocol identifier, and Frame Check Sequence. Flag: A single-byte delimiter

field set to 0x7E (binary value: 01111110) to indicate the boundary between the end of one PPP frame and the beginning of another PPP frame. Unlike SLIP, only a single Flag value appears between frames.

Protocol identifier: A two-byte field that identifies the upper-layer protocol ferried by the PPP frame.

Frame Check Sequence (FCS): A two-byte field that provides bit-level integrity checks for data as sent. (It's recomputed upon receipt, then compared to the sent value; if the two values agree, the assumption is that the data was transmitted successfully; if they disagree, the payload is discarded.)

Briefly discuss the following Ethernet II frame type fields and structures: Preamble, Destination Address Field, Source Address Field, and Type Field.
Preamble: The preamble is eight bytes long and consists of alternating 1s and 0s. As its name indicates, this special string of bits precedes the actual Ethernet frame itself, and is not counted as part of the overall frame length. The final byte ends in a pattern, the start frame delimiter (SFD), of 10101011, indicating the start of the Destination Address field. This field provides the necessary timing used by the receiver to interpret the 1s and 0s in a frame, and it builds in the time necessary for Ethernet circuitry to recognize and begin to read incoming data.

Destination Address Field: The Destination Address field is six bytes long and indicates the data link address (also referred to as the hardware address or MAC address) of the destination IP host. The destination address may be broadcast, multicast, or unicast. Address Resolution Protocol (ARP) is used to

obtain the hardware address of the destination IP host (if the destination is local), or the next-hop router (if the destination is remote).

Source Address Field: The Source Address field is six bytes long and indicates the sender's hardware address. This field can only contain a unicast address - it cannot contain a broadcast or multicast address.

Type Field: Two bytes long and identifies the protocol that is using this frame type.

Discuss the difference between the following ARP packet fields: Opcode field and the Protocol Type field. Opcode field: This field defines whether this ARP packet is a request or reply packet, and defines the type of address resolution taking place.

Protocol Type field: This field defines the protocol address type in use, and uses the standard protocol ID values that also are used in the Ethernet II frame structures.

The Target Hardware Address field indicates the desired target's hardware address, if known. In ARP requests, this field is typically filled with all 0s. In ARP replies, what should this field contain? The hardware address of the desired IP host if the sender and destination share a common data link.
or
The hardware address of the next router in the path to the destination if they do not share a common data link. This is known as the next-hop router to that IP host, in which that device will be the first of one or more routers that

will convey the data from sender to receiver. Each network, or router-to-router transition, is counted as a hop.

Discuss three ways in which a router entry can be placed in a routing table. The first way is through direct connection. For example, a router that is connected to networks 10. 1. 0. 0/16 and 10. 2. 0. 0/16 knows about both networks because its physical network interfaces reside on those subnets.

The second way is that it can be manually configured. To do this, you log on to the router and use the menus or command line to define a network that it can reach, the next hop, and any metrics. You repeat this process for every network you want to reach.

The third way that an entry can be placed in a routing table is dynamically, by using a routing protocol. Routers use routing protocols to share information about the various networks on an internetwork. Thus, you simply configure the protocol on each router, and the routers will convey Network Layer Reachability Information (NLRI) to each other.

How do link-state routing protocols differ from distance vector routing protocols? The first is that they do not route by rumor. Each router generates information about only its directly connected links, and these are passed around the entire network so every router in the area has an identical view of the network topology. Then the routers individually run an algorithm known as the Dijkstra algorithm to determine the optimal path through the internetwork.

The second major difference is that they do not periodically broadcast their entire tables. Link-state protocols build adjacencies with neighboring routers, and after an initial full exchange of information, they send only an update when a link state changes (for example, goes " up" or " down").

What is the purpose of the following ICMP message types: ICMP Redirect, ICMP Time Exceeded, and ICMP Parameter Problem? ICMP Redirect: Permits a gateway (router) on a nonoptimal route between sender and receiver to redirect traffic to a more optimal path.

ICMP Time Exceeded: Indicates that an IP datagram's TTL, or a fragmented IP datagram's reassembly timer, has expired; can indicate either a too-short TTL, or the presence of a routing loop on a network (which must be removed).

ICMP Parameter Problem: Indicates some error occurred while processing the IP header of an incoming datagram, causing that datagram to be discarded; catchall for ambiguous or miscellaneous errors, it indicates further investigation is required.

According to RFC 792, what is the relationship between IP and ICMP? ICMP provides a mechanism for gateways (routers) or destination hosts to communicate with source hosts.

ICMP messages take the form of specially formatted IP datagrams, with specific associated message types and codes.

ICMP is a required element in some implementations of TCP/IP, most notably those TCP/IP protocol stacks judged suitable for sale to the U. S. government, and ICMP is usually present to provide an essential part of IP's support fabric.

ICMP reports errors only about processing of non-ICMP IP datagrams. To prevent an endless loop of messages about error messages, ICMP conveys no messages about itself and provides information only about the first fragment in any sequence of fragmented datagrams.

What are the characteristics of the following packets: Windows 2008, Windows Vista, and Windows 7 Ping? The Identifier field is set to 512 decimal (or 0x200).
On the first echo sent, the Sequence Number field value is set to a multiple of 512 decimal (0x200). In each subsequent echo, this field is incremented by 256 decimal (02100).
The data field contains the value " abcdefghijklmnopqrstuvwabcdefghi."

Briefly define the following codes, currently assigned to the ICMP Destination Unreachable type number: Code 2: Protocol Unreachable, Code 3: Port Unreachable, and Code 5: Source Route Failed. Code 2: Protocol Unreachable: A host or router can send this error message to indicate that the protocol defined in the IP header cannot be processed.

Code 3: Port Unreachable: A host or router can send this reply to indicate that the sender does not support the process or application you are trying to reach.

Code 5: Source Route Failed: A router sends this ICMP reply to indicate that the router cannot use the strict or loose source routing path specified in the original packet.

How does inverse mapping determine live targets on a network? When a filtering device is detected between an attacker and his potential target, he can interrogate the routing device in an unusual way - he intentionally sends packets to vacant network addresses. Upon receipt of a packet destined for a nonexistent host, the intermediary router will gladly pass it on anyway (ICMP being a stateless protocol, the router knows no better). Once that packet reaches an internal router, however, one more knowledgeable in the valid and available network addresses, it will promptly reply with a Host Unreachable message for every bogus entry requested. The attacker then may logically deduce which addresses correspond to a live target.

What is firewalking? This is a two-phase attack method, involving an initial TRACEROUTE to discover hop count to a firewall appliance. Once this filtering device is identified by the TRACEROUTE, a second wave of attack follows, and this one consists of sending a packet with a TTL of one greater than the final hop count (between attacker and firewall). The goal is to elicit a Time Exceeded response from beyond the firewall, indicating a live and responsive target.

Describe some of the security issues for ICMPv6. ICMPv6 has built-in security features that are designed to prevent attacks sent from another network segment. These features include the value in the Hop Limit field being set at 255. Also, the source address of ICMPv6 packets must be either link-local or

unspecified (::/128) for all Router Advertisement and Neighbor Solicitation messages. However, no mechanism is currently specified that would prevent an attacker on the local network from exploiting ICMPv6 to compromise the network.

Authentication for ICMPv6 packet exchanges is managed using the IP Authentication Header (IPv6-AUTH) or the IP Encapsulating Security Payload Header (IPv6-ESP). IPv6-ESP also provides confidentiality for these exchanges.

ICMPv6 is protected by IPsec, but this presents a security bootstrap problem because IPsec is not available when a computer is at this state.

Briefly describe the following ICMPv6 message types: Router Solicitation, Router Advertisement, and Redirect. Router Solicitation (RS) (ICMPv6 type 133)—When an interface becomes active, a node may send a Router Solicitation message, asking any routers connected to the local link to identify themselves by sending their Router Advertisement messages immediately (rather than waiting for the next scheduled advertisement).

Router Advertisement (RA) (ICMPv6 type 134)—Routers periodically or upon request send out messages that contain at least one and possibly more of their own link-layer addresses, the network prefix for the local subnet, the Maximum Transmission Unit (MTU) for the local link, suggested hop limit values, and other parameters useful for nodes on the local link. Router Advertisement messages can also contain flagged parameters to indicate what type of address autoconfiguration process new nodes should use to join the network.

Redirect (ICMPv6 type 137)—When a router knows a better first-hop for a particular destination address (which could be off-link), it sends a Redirect message to the sender indicating that the sender should contact a different router to send subsequent packets.

Briefly describe the following ICMPv6 message types: Neighbor Solicitation and Neighbor Advertisement. Neighbor Solicitation (NS) (ICMPv6 type 135)—A node can send a Neighbor Solicitation message to find (or verify) the link-layer address for a local node, to see if that node is still available, or to check that its own address is not in use by another node, which is known as Duplicate Address Detection (DAD).

Neighbor Advertisement (NA) (ICMPv6 type 136)—When requested, or when its own link-layer address changes, a node sends a Neighbor Advertisement message that includes its IPv6 address and its link-layer address. This helps to establish physical adjacency (which is often more important than logical adjacency by address) to neighboring nodes

What information does an IPv6 Router Advertisement message contain for Ethernet interfaces? For an Ethernet interface, the Router Advertisement message is composed of the following: Ethernet header:

- Source address is the MAC address for the host interface.
- Destination address is 33: 33: 00: 00: 00: 01.

IPv6 header:

- Source address is the link-local address for the interface.
- Destination address is the link-local scope all-nodes multicast address FF02:: 1 or the source address for the interface.

Hop Limit:

- Set to 255 (an 8-bit integer value).

What information does an IPv6 Neighbor Solicitation message contain for Ethernet interfaces? For an Ethernet interface, the Neighbor Solicitation message is composed of the following:

Ethernet header:

- Source address is the MAC address of the host interface.
- Destination address is either the MAC address of the solicited-node address of the target (multicast NS) or the MAC address of the unicast address of the target (unicast NS).

IPv6 header:

- Source address is the IPv6 address of the interface or the unspecified address for DAD.
- Destination address is either the solicited-node address of the target (multicast NS) or the unicast address of the target (unicast NS).

Hop Limit:

- Set to 255 (an 8-bit integer value)

What is a conceptual host model? RFC 4861 does not exactly mandate how the ND process is to operate on all nodes; rather, it defines what must occur for the ND process to be successful. The ND definition for this operation is known as the conceptual host model, which represents information that a host should maintain, in some form, in order to communicate effectively in an IPv6 network. Some manufacturers have chosen different methods in their IPv6 protocol stacks to implement some of the component processes within ND to allow for proper communications with other IPv6-capable nodes.

The conceptual host model is primarily concerned with operational behavior by hosts. Routers have many of the same operational requirements, but they have additional needs such as routing operations controlled by routing protocols (if implemented) and some of the other data components that may be obtained and stored differently.

What neighbor data should be stored on a host? For each active network interface, a node needs to store all the following information:

Neighbor cache—A table of information containing the on-link address for each neighbor. It may include the link-layer address, the neighbor's state of reachability, and whether the neighbor is a host or a router.

Destination cache—A table of information containing data about destinations to which traffic has been sent, including both on-link and off-link nodes. The destination IPv6 address is mapped to the next-hop address of the neighbor. Data not related to ND may also be stored in the destination cache, such as the PMTU and round-trip timers. This list may also be updated from Redirect messages.

Prefix list—A table of information containing data from Router Advertisement messages of the on-link prefix addresses. In addition, each entry has an invalidation timer so it can expire prefixes as they become invalid. Link-local prefixes have an infinite invalidation timer regardless of whether a Router Advertisement message is received for the link-local prefix or not.

Default router list—This contains IP addresses of routers that have sent Router Advertisement messages. Each entry also includes its invalidation timer value.

Briefly describe the conceptual sending algorithm. For each active network interface, a node needs to For a node to communicate with a neighbor node, it needs to find out the IP address of the next-hop by examining its destination cache to learn the associated link-layer address by examining its neighbor cache. If the node does not have these addresses available, it invokes a process called " next-hop determination" to populate its caches and lists with its neighbor's addressing information. This process is known as the conceptual sending algorithm..

What are the processes involved with Neighbor Discovery? Address Resolution
Neighbor Unreachability Detection
Duplicate Address Detection
Router Discovery
Redirect Function

Provide a brief description of how DHCP works from a client perspective. Here's a brief rundown of how DHCP works, from a client perspective:

1. When TCP/IP is configured on the client computer, the Obtain an IP address automatically option is the only necessary set-up element. The DHCP service is automatic, which explains the terrific appeal that DHCP holds for network administrators and users alike.

2. The next time the workstation attempts to access the network (older versions of Windows must be rebooted first), it broadcasts a DHCP address request to the network because it has no IP address. It can make this broadcast because it is now configured as a DHCP client.

3. All DHCP servers present on the same broadcast domain receive the request and send back a message that indicates a willingness to grant an address lease, if an address is available.

4. The client accepts an address lease offer (usually the first one it receives) and sends a packet to the server that extended the offer.

5. In reply, the server proffers an IP address for a specific period of time (which is why it's called a lease) that the client uses thereafter.

6. When half the lease period expires, the client attempts to renew the lease. Usually, the DHCP server that granted the lease will renew it, but if it doesn't respond, the client tries to renew again at other times during the lease period. Only if the client is unable to renew its lease before expiration must that client repeat the DHCP request process, as described in Step 2.

What is the difference between the DHCP server software and the DHCP client software? DHCP client: The DHCP client software broadcasts requests for service and lease renewal requests on behalf of its clients, and it handles address and configuration data for the client when an address lease is granted. Windows Server 2008, Windows Vista, Windows 7, Macintosh, Linux, and UNIX machines all include built-in DHCP client software.

DHCP server: DHCP server software listens and responds to client and relay requests for address services. The DHCP server also manages address pools and related configuration data. Most current DHCP servers (UNIX and Windows Server 2008) can manage multiple address pools.

What is the difference between a manual address lease and a dynamic address lease? With a manual address lease, the administrator explicitly assigns an IP address manually by associating a client's hardware address with a specific IP address to be leased to that client. Use this type of address lease if you want DHCP to manage all IP addresses, but you also want to control some address assignments directly.

Dynamic: The DHCP server assigns addresses for specific periods of time. Use a dynamic address lease to assign addresses to clients or other machines when fixed IP addresses are not required. Given the prevalence of clients on most networks, this is the most prevalent type of DHCP address lease. Dynamic address leases come from the dynamic address pool, and represent those addresses not already reserved for static allocation.

How does DHCP integrate with DNS? Server addresses (and sometimes their associated services) are advertised using DNS, which resolves domain names into IP addresses, and vice versa.

DNS is not a dynamic environment, so all address updates must be entered manually (either through a GUI interface, on Windows Server 2008, or by editing text files on UNIX systems).

Client addresses usually come into play only when e-mail addresses of the formmust be resolved. E-mail servers can resolve this information from MX records associated with the client's domain name (not his or her IP address), so dynamic address resolution works perfectly well for clients and e-mail. Hence, client addresses typically have no impact on DNS, or vice versa, and can change, as needed.

List the four packets used by DHCP Discovery. DHCP Discover packet
DHCP Offer packet
DHCP Request packet
DHCP Acknowledgment packet

Briefly describe how to combine stateful and stateless address autoconfiguration. Routers on the local link can be configured to provide pointers to DHCPv6 servers that may provide only certain " other" types of network configuration information, such as DNS and time server addresses. The router in this case is configured with its RA A and L flags set to " on," the M flag set to " off," and the O flag set to " on." The router provides the network prefix, and the DHCPv6 server provides the DNS server information for hosts.

- Vector